

Proving linearisability via coarse-grained abstraction

Brijesh Dongol John Derrick

Department of Computer Science,
The University of Sheffield, S1 4DP,
United Kingdom

B.Dongol@sheffield.ac.uk, J.Derrick@dcs.shef.ac.uk

December 21, 2012

Abstract

Linearisability has become the standard safety criterion for concurrent data structures ensuring that the effect of a concrete operation takes place after the execution some atomic statement (often referred to as the linearisation point). Identification of linearisation points is a non-trivial task and it is even possible for an operation to be linearised by the execution of other concurrent operations. This paper presents a method for verifying linearisability that does not require identification of linearisation points in the concrete code. Instead, we show that the concrete program is a refinement of some coarse-grained abstraction. The linearisation points in the abstraction are straightforward to identify and the linearisability proof itself is simpler due to the coarse granularity of its atomic statements. The concrete fine-grained program is a refinement of the coarse-grained program, and hence is also linearisable because every behaviour of the concrete program is a possible behaviour its abstraction.

1 Introduction

With the increasing prevalence of concurrent computation in modern systems, development of concurrent data structures that enable a greater degree of parallelism have become increasingly important. To improve efficiency, programs that implement concurrent data structures often exhibit fine-grained atomicity and use atomic non-blocking compare-and-swap operations as their main synchronisation primitive. A consequence of these features is the increase in complexity of the programs, making their correctness harder to judge. Hence, formal verification of programs for concurrent data structures is known to be a necessity. There are even examples of errors being uncovered by formal verification in published algorithm that were previously believed to be correct [14].

The main correctness criterion for programs that implement concurrent data structures is linearisability [30], which allows one to view operations on concurrent objects as though they occur in some sequential order. Over the years, numerous approaches to verifying linearisability have been developed using a variety of different frameworks, and several of these approaches are partially/fully mechanised. Herlihy and Wing’s original paper use the notion of a possibilities mapping, which defines the set of possible abstract data structures that corresponds to each point of interleaving. Doherty et al [5, 15] use a simulation-based method using input/output automata with proofs mechanised using PVS. Vafeiadis et al use a framework that combines separation logic and rely/guarantee reasoning [46, 48]. An automated method based on this theory has been developed, but the method is known not to apply to a more complex programs [47]. Derrick et al have developed refinement-based methods that have been mechanised using the theorem prover KIV [11, 12, 13]. Turon and Wand propose a compositional rely/guarantee framework with separation logic to show that the concrete programs implement another so-called “obviously correct” program, which may or may not be linearisable [45]. Several other verification methods have been proposed which we discuss in more detail in Section 9.

Each of the existing methods above build on the fact that linearisability guarantees the existence of a so-called *linearisation point*, which is an atomic statement whose execution causes the effect of an operation to be felt.

“Linearisability provides the illusion that each operation applied by concurrent processes takes effect instantaneously at some point between its invocation and its response.” [30]

Hence, the methods in [5, 11, 12, 13, 15, 46, 47, 48] involve identification of linearisation points in the concrete code, and a proof that execution of a linearisation point does indeed correspond to an execution of the corresponding abstract operation. However, in many sophisticated programs, the linearisation points are not always immediately identifiable, and often require a high degree of expertise on the proof techniques as well as the program at hand. In more complex algorithms, it is even possible for some operations to be linearised by the execution of other concurrent operations and hence require the use of backwards reasoning techniques [5, 11, 13, 43, 46, 48]. Algorithms that require backwards reasoning are precisely those that cause difficulties for the method described in [47].

We present a method for verifying linearisability where we aim to establish a relationship between a fine-grained concrete program and a program in which the operations execute with coarse-grained atomicity. In particular, the fine-grained program is one that implements the one with coarse-grained atomicity. Groves [26] and separately Elmas et al [25] start with a coarse-grained program, which is incrementally refined to an implementation with finer-grained atomicity. Splitting the atomicity of a statement is justified using *reduction* [33], which ensures that the operations are immune to interleavings with other concurrent statements. However, because one must consider each pair of interleavings, reduction-based methods are not compositional, and hence do not scale well as the complexity of an operation increases.

We develop a framework that enables reasoning over the interval time in which a program executes, presenting an alternative to traditional reasoning over the pre/post states of a program and captures the possible interference that may occur during a program’s execution [18]. Our model incorporates fractional permissions [3] to simplify reasoning about conflicting accesses to shared variables [18]. Permissions are also used to model properties such as interference freedom and locality. Our framework also incorporates reasoning about pointer-based programs by allowing the domains of each state to be assumed to consist of variables and addresses, and take extra care when updating or evaluating pointers because the values at their addresses may be dynamically changing. The behaviour of a command is defined over an interval, and hence, for example, expression evaluation is assumed to take a number of steps. Note though that we do not take into account all the complexities of non-deterministic expression evaluation [28].

We develop interval-based theories of refinement, namely *behaviour refinement* (which is akin to operation refinement) and *data refinement* (which allows the state spaces of two programs to be linked using a simulation predicate). As far as we are aware, our formulation of data refinement over intervals is novel to this paper. Data refinement is used to link the abstract representation of the data structure (in which each operation takes place in a single atomic step) and the coarse-grained abstraction (in which operations execute on the same data representation as the implementation, but with a coarse-grained atomicity). Like Derrick et al [42, 13, 12], the data refinement proof encapsulates a proof that the coarse-grained program is linearisable with respect to the abstract representation. Then to show that the final implementation is linearisable, we show that it is a behaviour refinement of the coarse-grained abstraction. We use Treiber’s Stack [44] to illustrate our approach.

Interval-based methods for proving linearisability have also been proposed by Baumler et al [2]. However, their model assumes that a program executes with its environment by interleaving the statements of a program with those of its environment, as opposed to our model, that allows true concurrency, which allows one to model the inherent true parallelism in modern many/multicore systems. Furthermore, Baumler et al prove linearisability of the concrete program directly unlike our method in which we first show that the fine-grained (concrete) implementation refines a coarse-grained abstraction. Due to the coarse-granularity of the statements in the second program, its linearisation points are straightforward to identify, and the linearisation proof itself is simpler.

This paper is structured as follows. In Section 2, we present a formalisation of linearisability and an alternative definition that simplifies its proof. In Section 3, we present the Treiber stack, which we use as a running example throughout the paper. We present our interval-based framework in Section 4 and use it to define a semantics for a language that allows explicit control of a program’s atomicity (Section 5). We also develop a theory for refining the behaviour of commands within this framework. In Section 6, we develop the coarse-grained abstraction of the Treiber stack, present methods of verifying its linearisability using data refinement, and present the actual proof of its correctness. As part of this, we develop data refinement rules for our interval-based framework. We then develop methods for showing that a fine-grained program implements a coarse-grained abstraction. To this end, we develop compositional rely/guarantee-style rules in Section 7 and some high-level transformation rules specific to

CAS-based implementations in Section 7.3. We apply these to prove that the Treiber stack implements the coarse-grained abstraction in Section 8.

2 Linearisability

In this section, we present the original definition of linearisability [30] and an alternative definition that simplifies proofs of linearisability.

Two operations opi and opj of a concurrent program are said to execute *concurrently* iff the invocation of opi occurs after the invocation but before the response of opj . During a concurrent execution of two or more operations, the atomic statements of the operations may be arbitrarily interleaved. As a result, the effect of two concurrent operations may take place in any order and does not correspond to the ordering of invocations and responses. For example, Fig. 1 depicts a scenario where process r linearises before process p even though the invocation of p occurs before the invocation of r . Similarly, process q linearises after process q even though the response of process q occurs before the response of process p .

Not every ordering of invocations and responses is linearisable. In particular, linearisability requires that a chosen ordering of effects of the concurrent operations corresponds to a valid sequential history. For example, assuming that we start with an empty stack, history in Fig. 1 is linearisable whereas Fig. 2 is not [12]. In particular, history Fig. 1 can be linearised by selecting linearisation points marked by the crosses. In contrast, there is no possible selection of linearisation points for Fig. 2 that results in a valid sequential history because processes q and r both return x , even though there is only one concurrent $push(x)$ operation and execution started with an empty stack. We give a more formal presentation of these concepts in Example 2.2.

2.1 Herlihy and Wing's definition

To formalise linearisability using the nomenclature of Herlihy and Wing [30], we let $\text{seq}.X$ denote sequences of type X . We assume sequences start with index 0. An event is a tuple $Event \hat{=} Op \times Proc \times \{invoke, return\} \times \text{seq}.Val$, respectively corresponding to an operation identifier, a process identifier, the type of the event (invoke or return), and a sequence corresponding to the input/output parameters of the event. We use $op_p^I(k_1, k_2, \dots, k_n)$ and $op_p^R(k_1, k_2, \dots, k_n)$ to denote operations $(op, p, invoke, \langle k_1, \dots, k_n \rangle)$ and $(op, p, response, \langle k_1, \dots, k_n \rangle)$, respectively. Notations op_p^I and op_p^R denote an invoke and response events with no parameters, respectively. A *history* is a sequence of invocation of response events.

Example 2.1. For a stack data structure, sequences (1) and (2) below are possible histories of invocation and response events, where p, q and r are pairwise distinct processes.

$$\langle push_p^I(x), push_q^I(y), pop_r^I, push_q^R, push_p^R, pop_r^R(Empty) \rangle \quad (1)$$

$$\langle push_p^I(x), pop_q^I, pop_r^I, pop_q^R(x), push_r^R, pop_p^R(x) \rangle \quad (2)$$

A visualisation of histories (1) and (2) are given in Fig. 1 and Fig. 2, respectively. ♣

For $H \in \text{seq}.Event$ of invocations and responses, $H \upharpoonright p$ denotes the subsequence of H consisting of all invocations and responses of process p . Two histories H_1, H_2 are *equivalent* if for all processes p , $H_1 \upharpoonright p = H_2 \upharpoonright p$. An invocation $opi_p^I(x)$ *matches* a response $opj_q^R(y)$ iff $opi = opj$ and $p = q$. An invocation is *pending* in a history H iff there is no matching response to the invocation in H . We let $complete(H)$ denote the maximal subsequence of history H consisting of all invocations and matching responses in H , i.e., the history obtained by removing all pending invocations of H .

An operation op in a history is defined by an invocation $invocation.op$ followed by the next matching response $response.op$. For a history H , $<_H$ is an irreflexive partial order on operations where $opi <_H opj$ iff $response.opi$ occurs before $invocation.opj$ in H , i.e., opi and opj do not execute concurrently and opi occurs before opj .

Definition 2.1 (Sequential history). A history H is *sequential* iff the first element of H is an invocation and each invocation (except possibly the last) is immediately followed by its matching response.

Definition 2.2 (Linearisability [30]). A (concurrent) history HC is *linearisable* iff HC can be extended to a (concurrent) history HC' by adding zero or more matching responses to pending invocations such that $complete(HC')$ is equivalent to some sequential history HS and $<_{HC} \subseteq <_{HS}$.




Figure 1: History corresponding to (1)




Figure 2: History corresponding to (2)

Example 2.2. Using Definition 2.2, assuming that the stack is initially empty, history (1) is may be linearised by the following sequential history:

$$\langle pop_r^I, pop_r^R(Empty), push_q^I(y), push_q^R, push_p^I(x), push_p^R \rangle \quad (3)$$

Note that a single concurrent history may be linearised by several sequential histories. For example (1) can also be linearised by the following sequential history, in which case the order of the linearisation points of process p and q shown in Fig. 1 would be swapped.

$$\langle pop_r^I, pop_r^R(Empty), push_p^I(x), push_p^R, push_q^I(y), push_q^R \rangle \quad (4)$$

Linearising history (1) using (3) results in a final stack $\langle y, x \rangle$ with element y at the top, whereas (4) results in a final stack $\langle x, y \rangle$ with element x at the top.

Unlike (1), there is no valid linearisation of (2). ♣

The definition of linearisability allows histories to be extended with matching responses to pending invocations. This is necessary because some operations may be past their linearisation point, but not yet responded. For example consider the following stack history, where the stack is initially empty.

$$\langle push_p^I(x), pop_q^I, pop_q^R(x) \rangle \quad (5)$$

The effect of the invocation $push_p^I(x)$ has clearly been felt in (5) because the pop_q returns x , i.e., the linearisation point of $push_p^I(x)$ occurs before that of pop_q^I . We can validate this formally because (5) can be extended with a matching response to $push_p^I(x)$, then linearised by the following sequential history

$$\langle push_p^I(x), push_p^R, pop_q^I, pop_q^R(x) \rangle$$

2.2 An alternative definition of linearisability

Verifying linearisability (Definition 2.2) by reasoning at the level of histories of invocations and responses directly is clearly infeasible. Hence, we follow the methods of Derrick et al who link the concurrent and sequential histories using a matching function [42, 12, 13]. This then allows one to prove linearisability via data refinement. The key idea here is to distinguish between invocations that have and have not linearised but not yet returned. This allows one to determine whether or not a concrete program contributes to the abstract history.

One must first define matching pairs of events and pending invocation events within a sequence of events.

Definition 2.3 (Matching pair, Pending invocation). For a sequence of events H , we say $i, j \in \text{dom}.H$ forms a *matching pair* in H iff $mp_H(i, j)$ holds and i is a *pending invocation* in H iff $pending_H.i$ holds, where

$$\begin{aligned} mp_H.(i, j) &\hat{=} (i < j) \wedge (H.i.opid = H.j.opid) \wedge (H.i.proc = H.j.proc) \wedge \\ &\quad (H.i.type = \text{invoke}) \wedge (H.j.type = \text{response}) \\ &\quad (\forall k \bullet i < k < j \Rightarrow H.k.proc \neq H.i.proc) \\ pending_H.i &\hat{=} (H.i.type = \text{invoke}) \wedge (\forall j: \text{dom}.H \bullet j \geq i \Rightarrow \neg mp_H.(i, j)) \end{aligned}$$

Hence, two indices i and j form a matching pair in a sequence of events H iff j follows i , events $H.i$, $H.j$ are invocations and responses of the same operation by the same process and there are not other invocations/responses by process $H.i.proc$ between i and j . Index i of sequence H is pending if there is no matching index $j > i$.

Example 2.3. Given that H is the history corresponding to (1), $mp_H.(0, 4)$, $mp_H.(1, 3)$ and $mp_H.(2, 5)$ hold. If h is the history corresponding to (5), $mp_h.(1, 2)$ and $pending_h.0$ hold. ♣

Using these, one may now define the set of legal histories.

Definition 2.4 (Legal history). A history H is *legal* iff $legal.H$ holds, where

$$legal.H \quad \hat{=} \quad \forall i: \text{dom}.H \bullet \text{if}(H.i.type = \text{invoke}) \text{ then } (pending_H.i \vee \exists j: \text{dom}.H \bullet mp_H.(i, j)) \\ \text{else } (\exists j: \text{dom}.H \bullet mp_H.(j, i))$$

Hence, H is legal iff for every index corresponding to an invocation in H is either pending or has a matching index, and indices corresponding to responses have an earlier matching invocation. Using this, one may now define a notion of a lin-relation between two histories.

Definition 2.5 (Lin-relation). A history HC is said to be in a *lin-relation* with history HS with respect to a matching function f iff $linrel(HC, f, HS)$ holds, where

$$linrel(HC, f, HS) \quad \hat{=} \quad f \in \text{dom}.HC \Rightarrow \text{dom}.HS \wedge \quad (6)$$

$$(\forall i, j: \text{dom}.HC \bullet mp_{HC}.(i, j) \Rightarrow \{i, j\} \subseteq \text{dom}.f) \wedge \quad (7)$$

$$(\forall i: \text{dom}.HS \bullet HC.i = HS.(f.i)) \wedge \quad (8)$$

$$(\forall i, j: \text{dom}.HC \bullet i < j \wedge mp_{HC}.(i, j) \Rightarrow f.j = f.i + 1) \wedge \quad (9)$$

$$(\forall i, j, k, l: \text{dom}.HC \bullet j < k \wedge mp_{HC}.(i, j) \wedge mp_{HC}.(k, l) \Rightarrow f.j < f.k) \quad (10)$$

Hence, a concrete history HC is in a lin-relation with sequential history HS with respect to linearising function f iff f is a surjection (an onto function) mapping the indices of HC to the indices of HS (6), every pair of indices of HC that forms a matching pair in HC is in the domain of f (7), for every index i of the sequential history HC , element $HC.i$ is the same as element $HS.(f.i)$ (8), for every matching pair of indices i and j , $f.j$ is one greater than $f.i$ (9), and matching indices i, j must occur before matching indices k, l if j occurs before k (10).

We let $prefix.tt$ denote the set of all prefixes of a sequence tt and hence $ss \in prefix.tt$ denotes that sequence ss is a prefix of tt . Using the definitions above, we now present Derrick et al's alternative definition of linearisability [12].

Definition 2.6 (Linearisability). A (concurrent) history HC is *linearisable* with respect to a sequential history HS iff $linearisable(HC, HS)$ holds, where

$$linearisable(HC, HS) \quad \hat{=} \quad \exists HE: \text{seq}.Event \bullet HC \in prefix.HE \wedge legal.HE \wedge \exists f \bullet linrel(HE, f, HS)$$

Hence, concrete history HC is linearisable with respect to sequential history HS iff HC can be extended to a sequence of (return) events HE such that the extended history HE is legal and there exists a linearising function between HE and HS .

3 Example: The Treiber Stack

We present our methods via a verification of the Treiber Stack as a running example (see Fig. 3), which is a well-known program that implements a list-based concurrent stack [44]. Verification of this stack has become a standard exercise in the literature. The program uses a pointer **Top** of type **ptr_ctr** within which the pointer field **ptr** stores a pointer to the top of the stack and counter field **ctr** stores the number of times **Top** has been modified. Stack nodes have a field **key** for the value of the node and a next field **nxt**, which is a pointer to the next node of the stack.

Like many lock-free algorithms, the stack is implemented using an atomic non-blocking compare-and-swap $CAS(ae, \alpha, \beta)$ primitive, which takes an address-valued expression ae and variables α and β as input. If the value at ae is equal to α , the CAS updates the value at ae to β and returns *true*, otherwise (the value at ae is not equal to α) the CAS does not modify anything and returns *false*.

The stack processes may perform either execute a push **push(x)** or a pop operation **pop**. Within operation **push(x)**, the executing process uses local variables **n** and **t**. The process executing **push(x)** sets up a new node with value x (lines h_1 - h_2), then executes a try-CAS loop (lines h_3 - h_5), where the value of global variable **Top** is read and stored in local variable **t** (line h_3), **n.nxt** is set to be the local **t** (line

```

data node {key: Val; nxt: *node}
struct ptr_ctr {ptr: *node; ctr: nat}
var Top: *ptr_ctr
initially Top = (null, 0)

push(x)  $\hat{=}$ 
  h1: n := new(Node) ;
  h2: n.key := x ;
  repeat
    h3: t := *Top ;
    h4: n.nxt := t.ptr
    h5: until CAS(Top, t, (n, t.ctr + 1))

pop  $\hat{=}$ 
  repeat
    l1: t := *Top ;
    l2: if t.ptr = null then
      l3: rv := Empty ;
    l4: return
    else
      l5: tn := t.ptr → nxt ;
      l6: rv := t.ptr → key
    fi ;
    l7: until CAS(Top, t, (tn, t.ctr + 1)) ;
  l8: return rv

```

Figure 3: Push and pop operations of Treiber’s lock-free stack

h_4) and a $\text{CAS}(\text{Top}, t, (n, t.\text{ctr} + 1))$ is executed (line h_5). The loop terminates if the CAS is successful, otherwise lines h_3 - h_5 are re-executed. Operation **pop** stores **Top** in local variable **t** (line l_1). If **t** is null (line l_2) it returns empty (line l_3), otherwise it stores the value of $t.\text{ptr} \rightarrow \text{nxt}$ in local variable **tn** (line l_4) and the value of $t.\text{ptr} \rightarrow \text{key}$ in local variable **rv** (line l_5). The loop terminates if the CAS is successful (line l_6), otherwise it re-executes the loop body (i.e., from line l_1).

We verify this algorithm for an arbitrarily chosen set of processes P as follows, assuming that $AS(P)$, $LS(P)$, and $TS(P)$ denote, respectively, the abstract stack (which is implemented as a sequence), the coarse-grained abstraction (which is implemented as a linked list) and the fine-grained concrete implementation (i.e., the Treiber Stack). The approach we propose is to show that $TS(P)$ implements (i.e., is a refinement of) $LS(P)$, which itself implements $AS(P)$. Furthermore, we show that $LS(P)$ is *linearisable* with respect to $AS(P)$ [30, 12, 48, 46].

The idea is that the coarse-grained abstraction $LS(P)$ allows concurrency, but large parts of the code are sequential. This simplifies the proof of linearisability because a data refinement is performed from an abstract data structure representation in $AS(P)$ to $LS(P)$, as opposed to a data refinement from $AS(P)$ directly to the fine-grained program $TS(P)$. The data refinement between $AS(P)$ and $LS(P)$ is performed on an extended specification that includes histories of invocations and responses [12], which allows one to relate the refinement to Herlihy and Wing’s original definition. Verification that $TS(P)$ is a refinement of $LS(P)$ is performed using a number of decomposition and transformation theorems.

4 Interval-based framework

We start by presenting our programming syntax. We then present our interval-based framework, which has similarities to Interval Temporal Logic [37]. However, the underlying semantics consists of *complete streams* that map each time to a state, and adjoining intervals are assumed to be disjoint (adjoining intervals share a boundary in Interval Temporal Logic) [18]. To formalise the behaviour over an interval, we restrict one’s view of a stream to the interval under consideration. However, because streams encode the complete behaviour over all time, our framework allows properties outside the given interval to be considered in a straightforward manner [20, 22, 18].

4.1 Syntax

There are a number of well-established approaches to modelling program behaviour, e.g., Z , B , I/O automata. However, the use of these formalisms involve a non-trivial translation from the program to the model. We use a framework in which commands closely resemble program code, which simplifies the translation. The programs we consider often have a pointer-based structure and hence, as in separation

logic, we distinguish between variables and addresses in the domains of the program states [40].

We assume variable names are taken from the set Var , values have type Val , addresses have type $Addr \triangleq \mathbb{N}$, $Addr \cap Var = \emptyset$ and $Addr \subseteq Val$. A *state* over a set of locations $VA \in Var \cup Addr$ is a member of $State_{VA} \triangleq VA \rightarrow Val$ (i.e., a total function from VA to Val). A *state predicate* over VA is a function of type $State_{VA} \rightarrow \mathbb{B}$.

A location may correspond to a data type with a field identifiers of type $Field$. We assume that every data type with m number of fields is assigned m contiguous blocks of memory [46] and use $offset.f \in \mathbb{N}$ to return the offset of the field f . For example, for any node of the Treiber Stack, we have $offset.key = 0$ and $offset.next = 1$. We assume that expressions have the following syntax, where $k \in Val$ is a constant, $v \in Var$, ae is an address-valued expression, f is a field, e , e_1 and e_2 are expressions, \ominus is a unary operators and \oplus is a binary operator, respectively. Both \ominus and \oplus are abstractions of the possible unary and binary operators on expressions.

$$e ::= k \mid v \mid *ae \mid ae.f \mid ae \mapsto f \mid \ominus e \mid e_1 \oplus e_2$$

The semantics of expressions is given by function $eval$ which is defined below for a state σ , where ‘.’ denotes function application.

$$\begin{aligned} eval.k.\sigma &\triangleq k & eval.(ae.f).\sigma &\triangleq eval.ae.\sigma + offset.f \\ eval.v.\sigma &\triangleq \sigma.v & eval.(\ominus e).\sigma &\triangleq \ominus eval.(e.\sigma) \\ eval.(*ae).\sigma &\triangleq \sigma.(eval.ae.\sigma) & eval.(e_1 \oplus e_2).\sigma &\triangleq eval.(e_1.\sigma) \oplus eval.(e_2.\sigma) \end{aligned}$$

Hence, the evaluation of a constant (including an address) in any state is the constant itself and an evaluation of a variable is the value of the variable in the given state. Evaluation of $*ae$ returns the value of the state at the address that ae evaluates to, and $ae.f$ returns the address that ae evaluates to plus the offset of the field f . The interpretations over unary and binary operators are lifted in the normal manner. We assume that expressions are well-formed so that their evaluation in every state is always possible. We define a shorthand

$$ae \mapsto f \triangleq *(ae.f)$$

which returns the value of the state at address $ae.f$.

Definition 4.1 (Command syntax). For a state predicate c , variable v , expression e , address-valued expression ae , set of processes $P \subseteq Proc$, set of variables Z , and label l , the abstract syntax of a command is given by Cmd below, where $C, C_1, C_2, C_p \in Cmd$.

$$Cmd ::= Chaos \mid Idle \mid [c] \mid v := e \mid ae := e \mid C_1 ; C_2 \mid C_1 \sqcap C_2 \mid C^\omega \mid \parallel_{p:P} C_p \mid \llbracket Z/C \rrbracket \mid l:C \mid \text{INIT } c \bullet C$$

Thus, a command may be **Chaos** (which is chaotic and allows any behaviour), an **Idle** command, a guard $[c]$, assignments $v := e$ and $ae := e$, a sequential composition $(C_1 ; C_2)$, a non-deterministic choice $C_1 \sqcap C_2$, an iteration C^ω , a parallel composition $\parallel_{p:P} C_p$, a command C executing in a context Z , a command with a label l , or a command that executes from an initial state that satisfies c .

As an example, consider the simple sequential program in Fig. 4, which demonstrates use of the syntax in Definition 4.1. For the program in Fig. 4, we assume that S is a sequence of values.

Example 4.1. Operation $SPush(x)$ updates the sequence S by appending x to the start of S , $SEmpty(arv)$ models a pop operation that returns $Empty$, and $SDoPop(arv)$ models a pop operation on a non-empty stack, which checks to see if S is non-empty, sets the return value arv to be the value of $S.0$, and removes the first element by setting the new value of S to be the tail of S . A single stack operation is modelled by SPP , which consists of a non-deterministic choice between push operations that non-deterministically insert one of the possible key values onto the stack, or a pop operation. A possibly infinite number of stack operations starting from an empty stack is modelled by SS , which ensures that variable S is in the context of the program, that the initial value of S is $\langle \rangle$, and iterates over SPP . ♣

As a more complicated example, we now consider the program in Fig. 5, which is a formal model for the Treiber Stack from Fig. 3. We model a **ptr_ctr** a structure by type $Ptr_Ctr \triangleq Addr \times \mathbb{N}$. For $(pp, cc) \in Ptr_Ctr$, we define functions $ptr.(pp, cc) \triangleq pp$ and $ctr.(pp, cc) \triangleq cc$. The modification counter $ctr.(*Top)$ is used to avoid the ABA problem (see Example 4.3). Unlike a **ptr_ctr** structure, which is assumed to be accessed atomically, list nodes are objects whose fields may be accessed independently, and hence the *key* and *next* fields are modelled as having different addresses.

$$\begin{aligned}
SPush(x) &\hat{=} S := \langle x \rangle \wedge S \\
SEmpty(arv) &\hat{=} [S = \langle \rangle]; arv := Empty \\
SDoPop(arv) &\hat{=} [S \neq \langle \rangle]; arv := S.0; S := tail.S \\
SPop(arv) &\hat{=} SEmpty(arv) \sqcap SDoPop(arv) \\
SPP &\hat{=} (\prod_{x:Val} SPush(x)) \sqcap SPop(arv) \\
SS &\hat{=} \llbracket S / \text{INIT } S = \langle \rangle \bullet SPP^\omega \rrbracket
\end{aligned}$$

Figure 4: An abstract (sequential) stack specification

$$\begin{aligned}
Setup(p, x) &\hat{=} h_1: newNode(p, n_p); h_2: (n_p \cdot key) := x \\
TryPush(p) &\hat{=} h_3: t_p := *Top; h_4: (n_p \cdot next) := ptr.t_p; hf_5: CASFail_p(Top, t_p) \\
DoPush(p) &\hat{=} h_3: t_p := *Top; h_4: (n_p \cdot next) := ptr.t_p; ht_5: CASOK_p(Top, t_p, (n_p, ctr.t_p + 1)) \\
Push(p, x) &\hat{=} Setup(p, x); TryPush(p)^\omega; DoPush(p) \\
ToCAS(p) &\hat{=} l_1: t_p := *Top; lf_2: [ptr.t_p \neq null]; l_5: tn_p := ptr.t_p \mapsto next; l_6: rv_p := ptr.t_p \mapsto key \\
TryPop(p) &\hat{=} ToCAS(p); lf_7: CASFail_p(Top, t_p) \\
Empty(p, rv_p) &\hat{=} l_1: t_p := Top; lt_2: [ptr.t_p = null]; l_5: rv_p := Empty \\
DoPop(p, rv_p) &\hat{=} ToCAS(p); lt_7: CASOK_p(Top, t_p, (tn_p, ctr.t_p + 1)) \\
Pop(p, rv_p) &\hat{=} TryPop(p)^\omega; (Empty(p, rv_p) \sqcap DoPop(p, rv_p)) \\
TPP(p) &\hat{=} pidle: Idle; (\prod_{x:Val} Push(p, x)) \sqcap Pop(p, rv_p) \\
TInit &\hat{=} *Top = (null, 0) \wedge FAddr \subseteq Addr \setminus \{Top\} \\
TS(P) &\hat{=} \llbracket Top, FAddr / \text{INIT } TInit \bullet \parallel_{p:P} \llbracket t_p, n_p, tn_p, rv_p / TPP(p)^\omega \rrbracket \rrbracket
\end{aligned}$$

Figure 5: Formal model of the Treiber stack

Example 4.2. The (interval-based) semantics of our language is given in Section 5.1. The behaviours of both $CASOK_p$ and $CASFail_p$ are given in Example 5.1 and the behaviour of $newNode(p, n_p)$ is formalised in Example 5.2. Both commands require the use of permissions to control the atomicity (see Section 4.5).

The Treiber Stack consists global address $Top \in Addr$ and a set of free addresses $FAddr \subseteq Addr$. The initial value of Top is $(null, 0)$ and all addresses different from Top are free. We assume the existence of a garbage collector that gathers free pointers and returns them to $FAddr$.

Execution of $push(x)$ by process p is modelled by $Push(p, x)$. Within the push operation, we split the label h_5 into hf_5 and ht_5 to distinguish between execution of the failed and successful branches of the CAS, respectively. Commands $TryPush(p)$ and $DoPush(p)$ model executions of the loop body that fail and succeed in performing the CAS at h_5 , respectively. Within $Push(p, x)$, because we use an $^\omega$ iteration, command $TryPush(p)$ may be executed a finite (including zero) number of times, after which $DoPush(p)$ is executed. However, it is also possible for $TryPush(p)$ to be executed an infinite number of times in which case $DoPush(p)$ never executes. Such behaviour is allowable for the Treiber Stack, which only guarantees lock-freedom of its concurrent processes [7, 6, 16].

For the pop operation, we use $ToCAS(p)$ to model the statements executed by process p from the beginning of the loop up to the CAS at l_7 (via a failed test at l_2). As in $Push(p, x)$, label l_2 is split into lf_2 and lt_2 , and l_7 is split into lf_7 and lt_7 in the pop operation. The $TryPop(p)$ and $DoPop(p)$ commands model executions of the loop body that fail and succeed in executing the compare and swap at l_7 , respectively. Command $Empty(p)$ models an execution of the pop operation that returns empty. The $Pop(p)$ operation consists of a finite or infinite iteration of $TryPop(p)$ followed by an execution of either $Empty(p)$ or $DoPop(p)$.

The program is modelled as a parallel composition of processes, where each process repeatedly chooses either a pop or push operation non-deterministically, then executes the operation. \clubsuit

4.2 Intervals

A (discrete) *interval* (of type *Intv*) is a contiguous set of integers (of type $Time \hat{=} \mathbb{Z}$), i.e., we define

$$Intv \hat{=} \{ \Delta \subseteq Time \mid \forall t, t': \Delta \bullet \forall u: Time \bullet t \leq u \leq t' \Rightarrow u \in \Delta \}$$

We let $\text{lub}.\Delta$ and $\text{glb}.\Delta$ denote the *least upper* and *greatest lower* bounds of an interval Δ , respectively, and define $\text{lub}.\emptyset \hat{=} -\infty$ and $\text{glb}.\emptyset \hat{=} \infty$. If the size of Δ is infinite and $\text{glb}.\Delta \in \mathbb{Z}$, then $\text{lub}.\Delta = \infty$ (i.e., is not a member of \mathbb{Z}) and if Δ is infinite and $\text{lub}.\Delta \in \mathbb{Z}$ then $\text{glb}.\Delta = -\infty$. The *length* of a non-empty interval Δ is given by $\ell.\Delta \hat{=} \text{lub}.\Delta - \text{glb}.\Delta$, and we define the length of an empty interval \emptyset to be $\ell.\emptyset \hat{=} 0$. We define the following predicates on intervals.

$$\begin{aligned} \text{Inf}.\Delta &\hat{=} \text{lub}.\Delta = \infty \\ \text{Fin}.\Delta &\hat{=} \neg \text{Inf}.\Delta \\ \text{Empty}.\Delta &\hat{=} \Delta = \emptyset \end{aligned}$$

Hence, $\text{Inf}.\Delta$, $\text{Fin}.\Delta$ hold iff Δ has an infinite and finite least upper bound, respectively, and $\text{Empty}.\Delta$ holds iff Δ is empty.

We must often reason about two *adjoining* intervals, i.e., intervals that immediately precede or follow a given interval. For $\Delta, \Delta' \in Intv$, we define

$$\Delta \alpha \Delta' \hat{=} \Delta \neq \emptyset \wedge \Delta' \neq \emptyset \Rightarrow (\text{lub}.\Delta < \text{glb}.\Delta') \wedge (\Delta \cup \Delta' \in Intv)$$

Thus, $\Delta \alpha \Delta'$ holds if and only if Δ' immediately follows Δ and Δ and Δ' are disjoint. Furthermore, by conjunct $\Delta \cup \Delta' \in Intv$, the union of Δ and Δ' must be contiguous. Note that both $\Delta \alpha \emptyset$ and $\emptyset \alpha \Delta$ hold trivially.

4.3 Interval predicates

We aim to reason about the behaviours of a program over its interval of execution and hence define an interval-based semantics for the language in Definition 4.1. In particular, we define interval predicates [21, 22, 23, 18], which map an interval and a stream of states to a boolean. The stream describes the behaviour of a program over all time and an interval predicate describes the behaviour over the given interval.

A *stream* of behaviours over $VA \subseteq Var \cup Addr$ is given by the total function $Stream_{VA} \hat{=} Time \rightarrow State_{VA}$, which maps each possible time to a state over V and A . To reason about specific portions of a stream, we use *interval predicates*, which have type $IntvPred_{VA} \hat{=} Intv \rightarrow \mathcal{P}Stream_{VA}$. As with state expressions, we assume pointwise lifting of operators on stream and interval predicates. We assume pointwise lifting of operators on stream and interval predicates in the normal manner, e.g., if g_1 and g_2 are interval predicates, Δ is an interval and s is a stream, we have $(g_1 \wedge g_2).\Delta.s = (g_1.\Delta.s \wedge g_2.\Delta.s)$. When reasoning about properties of programs, we would like to state that whenever a property g_1 holds over any interval Δ and stream s , a property g_2 also holds over Δ and s . Hence, we define universal implication for $g_1, g_2 \in IntvPred$ as

$$g_1 \Rightarrow g_2 \hat{=} \forall \Delta: Intv, s: Stream \bullet g_1.\Delta.s \Rightarrow g_2.\Delta.s$$

We say $g_1 \equiv g_2$ holds iff both $g_1 \Rightarrow g_2$ and $g_2 \Rightarrow g_1$ hold.

We define two trivial interval predicates

$$\begin{aligned} \text{True}.\Delta.s &\hat{=} \text{true} \\ \text{False}.\Delta.s &\hat{=} \text{false} \end{aligned}$$

Like Interval Temporal Logic [37], for an interval predicate g , we say $(\Box g).\Delta.s$ holds iff g holds in each subinterval of Δ in stream s , say $(\Diamond g).\Delta.s$ holds iff g holds in some subinterval of Δ , and say $\text{prev}.g.\Delta.s$ holds iff there is some immediately preceding interval of Δ within which g holds in s . More formally, we define:

$$\begin{aligned} (\Box g).\Delta.s &\hat{=} \forall \Delta': Intv \bullet \Delta' \subseteq \Delta \Rightarrow g.\Delta'.s \\ (\Diamond g).\Delta.s &\hat{=} \exists \Delta': Intv \bullet \Delta' \subseteq \Delta \wedge g.\Delta'.s \\ \text{prev}.g.\Delta.s &\hat{=} \exists \Delta' \bullet \Delta' \alpha \Delta \wedge g.\Delta'.s \end{aligned}$$

The *chop* operator ‘;’ is a basic operator, where $(g_1 ; g_2).\Delta$ holds iff either interval Δ may be split into two parts so that g_1 holds in the first and g_2 holds in the second, or the least upper bound of Δ is ∞ and g_1 holds in Δ . Thus, we define

$$(g_1 ; g_2).\Delta.s \quad \hat{=} \quad (\exists \Delta_1, \Delta_2 : Intv \bullet (\Delta = \Delta_1 \cup \Delta_2) \wedge (\Delta_1 \propto \Delta_2) \wedge g_1.\Delta_1.s \wedge g_2.\Delta_2.s) \vee (\text{lub}.\Delta = \infty \wedge g_1.\Delta.s)$$

Note that Δ_1 may be empty in which case $\Delta_2 = \Delta$, and similarly if Δ_2 is empty then $\Delta_1 = \Delta$. In the definition of chop, we allow the second disjunct $\text{lub}.\Delta = \infty \wedge g_1.\Delta$ to allow for g_1 to model an infinite (divergent or non-terminating) program.

We define the possibly infinite iteration of an interval predicate p as follows. We assume that interval predicates are ordered using universal implication ‘ \Rightarrow ’ and that the stream and interval are implicit on both sides of the definition.

$$g^\omega \quad \hat{=} \quad \nu z \bullet (g ; z) \vee \text{Empty}$$

Thus, g^ω is a greatest fixed point that defines either finite or infinite iteration of g [24].

Properties that hold over a larger interval may be decomposed into properties of the subintervals if the interval predicate that formalises the property splits and/or joins [27, 22]. We also find it useful to reason about properties that widen, where a property holds over a larger interval if it holds over any subinterval.

Definition 4.2 (Splits, Joins, Widens). Suppose g is an interval predicate. We say

- g *splits* iff $g \Rightarrow \Box g$, i.e., if g holds over an interval Δ , then g must hold over all subintervals of Δ ,
- g *joins* iff $(g ; g^\omega) \Rightarrow g$, i.e., for any interval Δ if g iterates over Δ , then g must hold in Δ , and
- g *widens* iff $\Diamond g \Rightarrow g$, i.e., if g holds over some subinterval of Δ , then g holds over the interval Δ .

4.4 Evaluating state predicates over intervals

The values of an expression e at the left and right ends of an interval Δ with respect to a stream s are given by $\overleftarrow{e}.\Delta.s$ and $\overrightarrow{e}.\Delta.s$, respectively, which are defined as

$$\begin{aligned} \overleftarrow{e}.\Delta.s &\quad \hat{=} \quad eval.e.(s.(glb.\Delta)) \\ \overrightarrow{e}.\Delta.s &\quad \hat{=} \quad eval.e.(s.(lub.\Delta)) \end{aligned}$$

Note that $\overleftarrow{e}.\Delta.s$ is undefined if $glb.\Delta = -\infty$ and $\overrightarrow{e}.\Delta.s$ is undefined if $inf.\Delta$. Similarly, we use the following notation to denote that c holds at the beginning and end of the given interval Δ with respect to a stream s , respectively.

$$\begin{aligned} \overleftarrow{c}.\Delta.s &\quad \hat{=} \quad glb.\Delta \notin \{-\infty, \infty\} \wedge c.(s.(glb.\Delta)) \\ \overrightarrow{c}.\Delta.s &\quad \hat{=} \quad lub.\Delta \notin \{-\infty, \infty\} \wedge c.(s.(lub.\Delta)) \end{aligned}$$

It is often useful to specify that a state predicate holds on a point interval. Hence we define

$$[c].\Delta.s \quad \hat{=} \quad \exists t : Time \bullet \Delta = \{t\} \wedge c.(s.t)$$

Two useful operators for evaluating state predicates over an interval are $\Box c$ and $\Diamond c$, which ensure that c holds in *all* and *some* state of the given stream within the given interval, respectively. Thus, for an interval Δ and stream s , we define:

$$\begin{aligned} (\Box c).\Delta.s &\quad \hat{=} \quad \forall t : \Delta \bullet c.(s.t) \\ (\Diamond c).\Delta.s &\quad \hat{=} \quad \exists t : \Delta \bullet c.(s.t) \end{aligned}$$

As demonstrated using Example 4.3 below, operator \Diamond may be used to give a straightforward formalisation of the ABA problem.

Example 4.3 (The ABA problem). To effectively implement CAS-based operations, where $\text{CAS}(\text{ae}, \alpha, \beta)$ is executed by a process p , operations are often structured so that the value at address ae is stored as value of variable α , some processing is performed on α and the result stored in a variable β . Then a $\text{CAS}(\text{ae}, \alpha, \beta)$ is executed to attempt updating the value at ae to β . If the CAS fails, then the environment must have modified the value at ae since it was last read as α . However, it is possible for a CAS to succeed (when it should have failed) if the environment exhibits so-called ABA-like behaviour [5], where the value at ae changes from say α to β and then back to α . ABA-like behaviour of an expression e is formalised using the following interval predicate:

$$\text{ABA}.e \hat{=} \exists k_1, k_2 \in \text{Val} \bullet k_1 \neq k_2 \wedge (\diamond(e = k_1); \diamond(e = k_2); \diamond(e = k_1))$$

Hence $(\text{ABA}.e).\Delta$ holds iff Δ can be partitioned into three adjoining intervals Δ_1 , Δ_2 and Δ_3 such that the value at e is k_1 sometime within Δ_1 , then changes to k_2 sometime within Δ_2 and back to k_1 sometime within Δ_3 . Note that $\text{ABA}.e$ allows the value of e to change several times when changing from k_1 to k_2 then to k_1 . ♣

We say a location va is *stable* at time t in stream s (denoted $\text{stable_at}.va.t.s$) iff the value of va in s at time t does not change from its value at time $t - 1$, i.e.,

$$\text{stable_at}.va.t.s \hat{=} s.t.va = s.(t-1).va$$

Variable va is *stable* over an interval Δ in a stream s (denoted $\text{stable}.va.\Delta.s$) iff the value of va is stable at each time within Δ . A set of locations VA is *stable* in Δ (denoted $\text{stable}.VA.\Delta$) iff each variable in VA is stable in Δ . Thus, we define:

$$\begin{aligned} \text{stable}.va.\Delta.s &\hat{=} \forall t: \Delta \bullet \text{stable_at}.va.t.s \\ \text{stable}.VA.\Delta &\hat{=} \forall va: VA \bullet \text{stable}.va.\Delta \end{aligned}$$

Note that every location is stable in an empty interval and the empty set of locations is stable in any interval, i.e., both $(\text{stable}.VA).\emptyset$ and $\text{stable}.\emptyset.\Delta$ hold trivially.

4.5 Read/write permissions and interference

As we shall see in Section 5.1, the behaviour a process executing a command is formalised by an interval predicate, and the behaviour of a parallel execution over an interval is given by the conjunction of these behaviours over the same interval. Because the state-spaces of the two processes are potentially overlapping, there is a possibility that a process writing to a variable conflicts with a read or write to the same variable by another process. To ensure that such conflicts do not take place, we follow Boyland's idea of mapping variables to a *fractional permission* [3], which is *rational* number between 0 and 1. A process has write-only access to a variable v if its permission to access v is 1, has read-only access to v if its permission to access v is above 0 but below 1, and has no access to v if its permission to access v is 0. Note that we restrict access so that a process may not have both read and write permission to a variable. Because a permission is a rational number, read access to a variable may be split arbitrarily (including infinitely) among the processes of the system. However, at most one process may have write permission to a variable in any given state. Note that the precise value of the read permission is not important, i.e., there is no notion of priority among processes based on the values of their read permissions.

We assume that every state contains a *permission* variable Π whose value in state $\sigma \in \text{State}_V$ is a function of type

$$V \rightarrow \text{Proc} \rightarrow \{n: \mathbb{Q} \mid 0 \leq n \leq 1\}$$

Note that it is possible for permissions to be distributed differently within states σ , σ' even if the values of the standard variables in σ and σ' are identical, i.e., it is possible to get $\sigma.\Pi \neq \sigma'.\Pi'$ even if $(\{\Pi\} \triangleleft \sigma) = (\{\Pi\} \triangleleft \sigma')$ holds, where ' \triangleleft ' denotes the domain anti-restriction.

Definition 4.3 (Permission). A process $p \in \text{Proc}$ has *write-permission* to variable va in state σ iff $\sigma.\Pi.va.p = 1$, has *read-permission* to va in σ iff $0 < \sigma.\Pi.va.p < 1$, and has *no-permission* to access va in σ iff $\sigma.\Pi.va.p = 0$ holds.

We introduce the following shorthands, which define state predicates for a process p to have read-only and write-only permissions to a variable va , and to be denied permission to access va .

$$\mathcal{R}.va.p \hat{=} 0 < \Pi.va.p < 1 \quad \mathcal{W}.va.p \hat{=} \Pi.va.p = 1 \quad \mathcal{D}.va.p \hat{=} \Pi.va.p = 0$$

In the context of a stream s , for any time $t \in \mathbb{Z}$, process p may only write to and read from va in the transition step from $s.(t-1)$ to $s.t$ if $\mathcal{W}.va.p$ and $\mathcal{R}.va.p$, respectively. Thus, $\mathcal{W}.va.p$ does not grant process p permission to write to va in the transition from $s.t$ to $s.(t+1)$ (and similarly $\mathcal{R}.va.p$).

It is straightforward to use fractional permissions to characterise interference within a set of processes. For $P \subseteq \text{Proc}$ and $VA \subseteq \text{Var} \cup \text{Addr}$, we define

$$\mathcal{I}.VA.P \hat{=} \exists va: VA, p: \text{Proc} \setminus P \bullet \mathcal{W}.va.p$$

which states that there may be interference on locations in VA from a process different from those in P . We use $\mathcal{I}.VA.p$ to denote $\mathcal{I}.VA.\{p\}$ for a singleton set $\{p\}$ and $\mathcal{I}.va.P$ to denote $\mathcal{I}.\{va\}.P$ for a singleton set $\{va\}$. This characterisation of interference is particularly useful in this paper where we use rely conditions (see Section 7.1) to formalise the behaviour of the environment. For example, to state that environment of p does not modify a variable v , we simply ensure that the rely condition implies $\Box \neg \mathcal{I}.v.p$.

Such notions are particularly useful because we aim to develop rely/guarantee-style reasoning, where we use rely conditions to characterise the behaviour of the environment. To state an assumption that there is no interference on va during the execution of a command that uses va , one may introduce $\Box \neg \mathcal{I}.va.p$ as a rely condition to the command (see Example 7.1).

We define some conditions on streams using fractional permissions that formalise our underlying assumptions on access permissions.

HC1 If no process has write access to $va \in \text{Var} \cup \text{Addr}$ within an interval, then the value of va does not change within the interval, i.e.,

$$\Box(\forall p: \text{Proc} \bullet \neg \mathcal{W}.va.p) \Rightarrow \text{stable}.va$$

HC2 The sum of the permissions of the processes on any location va is at most 1, i.e.,

$$\Box((\sum_{p \in \text{Proc}} \Pi.va.p) \leq 1)$$

For the rest of this paper, we implicitly assume that these conditions hold. Alternatively, one could make the conditions explicit by adding them to the rely conditions of the programs under consideration.

Using these healthiness conditions, we obtain a number of relationships between the values of a variable and the permissions that a process has to access the variable. For example, we may prove that if a process has read permission to a variable, then no process has write permission to the variable. Furthermore, if over an interval no process has write permission to a variable, then the variable must be stable over the interval.

Lemma 1. *Both of the following hold for any location $va \in \text{Var} \cup \text{Addr}$*

$$\Box((\exists p: \text{Proc} \bullet \mathcal{R}.va.p) \Rightarrow (\forall p: \text{Proc} \bullet \neg \mathcal{W}.va.p)) \quad (11)$$

$$\Box(\forall x: \text{Proc} \bullet \neg \mathcal{W}.va.p) \Rightarrow \text{stable}.va \quad (12)$$

One may also define additional properties. For instance, a set of locations VA may be declared to be local within a set of processes P using the following interval predicate.

$$\text{Local}.VA.P \hat{=} \forall va: VA \bullet \Box \left((\forall q: \text{Proc} \setminus P \bullet \mathcal{D}.va.q) \wedge ((\forall p: P \bullet \neg \mathcal{W}.v.p) \Rightarrow (\forall p: P \bullet \mathcal{R}.v.p)) \right)$$

Hence, if $\text{Local}.VA.P$, then no process outside of P has permission to access the locations in VA and furthermore, if no process in P has write permission to a location va in VA , then all processes in P automatically obtain read permission to va .

It is important to be able to determine the set of all variables and addresses that must be accessed in order to evaluate an expression in a given state. The set of variables that are accessed is state dependent

because an expression may involve pointers and the address that the pointer points to may vary with the state. Hence, we define a function *accessed*, which returns the set of locations (i.e., variables and addresses) accessed. Below, we assume that $k \in Val$, $v \in Var$, $a \in Addr$, $f \in Field$, $t \in Time$, e_1 , and e_2 are expressions and s is a stream.

$$\begin{array}{ll}
accessed.k.\sigma & \hat{=} \{\} \\
accessed.v.\sigma & \hat{=} \{v\} \\
accessed.(*ae).\sigma & \hat{=} \{eval.ae.\sigma\} \cup accessed.ae.\sigma
\end{array}
\qquad
\begin{array}{ll}
accessed.(ae.f).\sigma & \hat{=} accessed.ae.\sigma \\
accessed.(\oplus e).\sigma & \hat{=} accessed.e.\sigma \\
accessed.(e_1 \oplus e_2).\sigma & \hat{=} accessed.e_1.\sigma \cup accessed.e_2.\sigma
\end{array}$$

We define the following predicate, which states that process p has permission to read each of the locations required to evaluate expression e in state σ .

$$ReadAllLocs.e.p.\sigma \hat{=} \forall va: accessed.e.\sigma \bullet \mathcal{R}.va.p$$

The following interval predicates state that process p writes to a location of e within interval Δ , that there is no interference on e within Δ , and that all other processes are denied access to the locations in e .

$$\begin{array}{ll}
(WriteSomeLoc.e.p).\Delta.s & \hat{=} \Diamond(\exists va \bullet va \in accessed.e \wedge \mathcal{W}.va.p) \\
(IntFree.e.p).\Delta.s & \hat{=} \Box(\forall va \bullet va \in accessed.e \Rightarrow \neg \mathcal{I}.va.p) \\
(OnlyAccessedBy.e.p).\Delta.s & \hat{=} \Box(\forall va \bullet \forall q: Proc \setminus \{p\} \bullet va \in accessed.e \Rightarrow \mathcal{D}.va.q)
\end{array}$$

5 Interval-based semantics of parallel programs

In Section 5.1, we present our interval-based semantics of the programming model, and in Section 5.2 we present the concept of enforced properties, which enable a behaviour of a command to be constrained. We present a theory for refining program behaviour in Section 5.3.

5.1 Semantics of commands

We use the following interval predicates to formalise the semantics of the commands in Definition 4.1, where p is a process, $va \in Var \cup Addr$ is a location and $Z \subseteq Var \cup Addr$ is a set of locations, e is an expression, k a constant and c is a state predicate.

$$\begin{array}{ll}
idle_{p,Z} & \hat{=} \forall va: Z \bullet \Box \neg \mathcal{W}.va.p \\
eval_{p,Z}(e, k) & \hat{=} \Diamond(e = k \wedge ReadAllLocs.e.p) \wedge idle_{p,Z} \\
update_{p,Z}(va, k) & \hat{=} \begin{cases} idle_{p,Z \setminus \{va\}} \wedge \neg Empty \wedge \Box(va = k \wedge \mathcal{W}_p.va) & \text{if } va \in Var \\ idle_{p,Z \setminus \{va\}} \wedge \neg Empty \wedge \Box(*va = k \wedge \mathcal{W}_p.va) & \text{if } va \in Addr \end{cases}
\end{array}$$

Hence, $idle_{p,Z}$ states that process p does not write to any of the locations in Z within the given interval. Interval predicate $eval_{p,Z}(e, k)$ models the evaluation of expression e to a value k by process p in context Z , where $eval_{p,Z}(e, k).\Delta.s$ holds iff there is a state $s.t$ (for $t \in \Delta$) such that the value of e in state $s.t$ is k and p can read each of the locations needed to evaluate e in $s.t$. Furthermore, p does not write to any of the variables of the context Z within Δ . Interval predicate $update_{p,Z}(va, k)$ models the modification of location va to value k by process p executing in a context Z , where $update_{p,Z}(va, k).\Delta.s$ holds iff within s , throughout Δ , the value of va is k , p has write permission to va and does not have write permission to any other variable in Z . Additionally, $\neg Empty$ holds to ensure that va is actually updated — this is necessary because $\Box c$ trivially holds for an empty interval.

We obtain the following lemma, which relates write permissions to expression evaluation and stability of a variable.

Lemma 2. *Suppose $V, Z \subseteq Var$, $p \in Proc$, e is an expression and $k \in Val$ is a constant. Then each of the following hold:*

$$idle_{p,Z} ; eval_{p,Z}(e, k) \Rightarrow eval_{p,Z}(e, k) \tag{13}$$

$$eval_{p,Z}(e, k) ; idle_{p,Z} \Rightarrow eval_{p,Z}(e, k) \tag{14}$$

$beh_{p,Z}.Chaos \hat{=} True$	$beh_{P,Z}.(C_1 ; C_2) \hat{=} beh_{P,Z}.C_1 ; beh_{P,Z}.C_2$
$beh_{p,Z}.Idle \hat{=} idle_{p,Z}$	$beh_{P,Z}.(C_1 \sqcap C_2) \hat{=} beh_{P,Z}.C_1 \vee beh_{P,Z}.C_2$
$beh_{p,Z}.[c] \hat{=} eval_{p,Z}.(c, true)$	$beh_{p,Z}.C^\omega \hat{=} (beh_{p,Z}.C)^\omega$
$beh_{p,Z}.(\text{INIT } c \bullet C) \hat{=} prev.\vec{c} \wedge beh_{p,Z}.C$	$beh_{p,Z}.(l : C) \hat{=} \Box(pc_p = l) \wedge beh_{p,Z}.C$
$beh_{p,Z}.(v := e) \hat{=} \exists k \bullet eval_{p,Z}(e, k) ; update_{p,Z}(v, k)$	
$beh_{p,Z}.(ae := e) \hat{=} \exists k, a \bullet eval_{p,Z}(ae, a) \wedge eval_{p,Z}(e, k) ; update_{p,Z}(a, k)$	
$beh_{P,Z}.(\parallel_{p:P} C_p) \hat{=} \begin{cases} True & \text{if } P = \emptyset \\ beh_{p,Z}.C_p & \text{if } P = \{p\} \\ \exists P_1, P_2 \bullet (P_1 \cup P_2 = P) \wedge (P_1 \cap P_2 = \emptyset) \wedge \\ \quad beh_{P_1,Z}.(\parallel_{p:P_1} C_p) ; Idle \wedge beh_{P_2,Z}.(\parallel_{p:P_2} C_p) ; Idle & \text{otherwise} \end{cases}$	
$beh_{P,Z}.[Y/C] \hat{=} Local.P.Y \wedge (Z \cap Y = \emptyset) \wedge beh_{P,Z \cup Y}.C$	

Figure 6: Formalisation of behaviour function

Definition 5.1 (Behaviour). The *behaviour* of a command C given by the abstract syntax in Definition 4.1 executed by a non-empty set of processes P in a context $Z \subseteq Var \cup Addr$ is given by interval predicate $beh_{P,Z}.C$, which is defined inductively in Fig. 6.

Within $beh_{P,Z}$ the context Z defines the set of locations that the processes in P may or may not modify. For example, command `Idle` executed by process p in a context Z should not write to the locations in Z . Similarly, an assignment $v := e$ should not write to locations in $Z \setminus \{v\}$.

In the description below, we assume that the executing process is p and the command under consideration occurs within the scope of a context Z . The behaviour of command `Idle` states that p does not write to any of the locations in Z . The behaviour of $[c]$ states that c evaluates to true in some state within the interval, that process p has permission to read the locations of c in this state and, and that p does not write to any of the locations in Z . For example, within $ToCAS(p)$ (Fig. 5), $[ptr.t_p \neq null]$ holds iff there is a state in which p can read t_p and the value of $ptr.t_p$ is non-null.

Execution of a variable assignment $v := e$ consists of two parts, where the value of e is evaluated to k , and then the value of v is updated to k . The executing process must have permission to read the locations of e within during the evaluation, and must have permission to write to v during the update. Note that because we assume true concurrency, the evaluation is non-deterministic (in the sense of [18, 28]). For example, consider the assignment $t_p := *Top$ within $ToCAS(p)$ (Fig. 5). It is possible for $*Top$ to change multiple times within the interval in which $*Top$ is evaluated due to the execution of other processes. The value returned by the evaluation of $*Top$ depends on the time at which the value at Top is read. Because t_p is local to p , the update to t_p may always be executed. Assignment $ae := e$ is similar to a variable assignment, but the command must additionally must evaluate the address-valued expression ae to determine the address to be updated.

The behaviours of sequential composition, non-deterministic assignment and iteration are modelled in a straightforward manner using chop, disjunction and iteration of interval predicates. The parallel composition is chaotic if the given set of processes is empty and behaves as a single process if the set of processes is singleton. Otherwise the given set processes is partitioned into two disjoint subsets and the behaviours of both subsets occur in the given interval. Note that the two or more parallel processes may access the same shared variables — the manner in which these variables are accessed is controlled using fractional permissions. We allow `Idle` to be executed within the parallel composition to allow asynchronous termination [18].

The behaviour of $[Y/C]$ is the behaviour of C in an extended context Y no other process is given permission to access locations in Y during the execution of C . The behaviour of a labelled command assumes the existence of an auxiliary program counter variable pc_p local to each process p . Execution of $l : C$ by process p guarantees that pc_p has value l throughout the interval of execution. Unlike [10, 13, 15, 17] where labels strictly correspond to the atomic portions of the programs under consideration, we only use labels to determine auxiliary information and the same label may correspond to a number of atomic steps.

5.2 Enforced conditions

We introduce a construct for defining an enforced condition, which restricts the behaviour of a command so that the property being enforced is guaranteed to hold [17].

Definition 5.2 (Enforced condition). For interval predicate d and command C , we let $\text{ENF } d \bullet C$ denote a command with an *enforced condition* d , where

$$\text{beh}_{P,Z}(\text{ENF } d \bullet C) \triangleq d \wedge \text{beh}_{P,Z}.C$$

Hence, $\text{ENF } d \bullet C$ executes as C and in addition guarantees that d holds [17, 22]. Note that if $\neg d$ holds, then $\text{ENF } d \bullet C$ has no behaviours. Enforced conditions may be used to state properties of an implementation that may not be easily expressible as a command. We present three examples relevant to the Treiber stack to illustrate the use of enforced properties. Although each of the examples only uses enforced properties on fractional permissions, the theory of enforced properties is more general as it allows any interval predicate to be enforced [22, 21, 18].

Example 5.1 (Compare-and-swap). A $\text{CAS}(\text{ae}, \alpha, \beta)$ is atomic on ae , i.e., the value at address ae is guaranteed not to be modified by the environment over the interval in which $\text{CAS}(\text{ae}, \alpha, \beta)$ is executed. Enforced properties, state predicate evaluation and fractional permissions are used together to formalise the behaviour of a $\text{CAS}(\text{ae}, \alpha, \beta)$ executed by a process p , as follows:

$$\begin{aligned} \text{CASOK}_p(\text{ae}, \alpha, \beta) &\triangleq \text{ENF } \text{OnlyAccessedBy}.(*\text{ae}).p \bullet [* \text{ae} = \alpha] ; \text{ae} := \beta \\ \text{CASFail}_p(\text{ae}, \alpha) &\triangleq [* \text{ae} \neq \alpha] \\ \text{CAS}_p(\text{ae}, \alpha, \beta) &\triangleq \text{CASOK}_p(\text{ae}, \alpha, \beta) \sqcap \text{CASFail}_p(\text{ae}, \alpha) \end{aligned}$$

Note that due to the enforced properties $\text{OnlyAccessedBy}. \text{ae}.p$ within CASOK_p , the environment is denied access to the locations needed to evaluate ae . Although the CAS may take a number of atomic steps to execute, access to the test and set of ae occurs without any process accessing the locations in ae . A CAS that performs the test and set in a single transition may be considered to be an implementation of this specification.

Because the CAS may take multiple steps, it is possible for the values of α and β to change. To achieve more deterministic behaviour over the interval, (in the sense of [28]), α and β are typically local variables of the executing process, and hence, by **HC4**, α and β cannot be modified by the environment of p . A CAS that performs a test and set in a single transition typically places further restrictions on the structure of ae to ensure implementability. ♣

Example 5.2 (New nodes). We assume that the set of free addresses is given by $F\text{Addr}$. To ensure that two different processes are not assigned the same free address, we must ensure that there is no interference to $F\text{Addr}$ while p accesses $F\text{Addr}$. This may be achieved via an enforced condition on $F\text{Addr}$.

$$\begin{aligned} \text{newNode}(p, v) &\triangleq \text{ENF } \text{OnlyAccessedBy}.F\text{Addr}.p \bullet \\ &\quad \bigcap_{fn \in F\text{Addr}} [fn + \text{offset}.nxt \in F\text{Addr}] ; v := fn ; F\text{Addr} := F\text{Addr} \setminus \{fn, fn + \text{offset}.nxt\} \end{aligned}$$

Hence, $\text{newNode}(p, v)$ ensures that all processes different from p are denied access to the current set of free nodes $F\text{Addr}$. Furthermore, it non-deterministically chooses a free node from fn such that $fn + 1$ is a free node, then assigns fn to v , and removes both fn and $fn + \text{offset}.nxt$ (i.e., fn and $fn + 1$) from the set of available free locations. ♣

Example 5.3. To further illustrate the use of permissions and enforced properties, consider the program in Fig. 7, which extends the abstract program in Fig. 4 by allowing the operations to be executed concurrently by the processes in $P \subseteq \text{Proc}$. However, due to the enforced permissions, each process is guaranteed to update the stack without interference. Lynch [35] refers to such a program as a *canonical specification* of the concurrent stack. Note that some idling must be allowed both before and after the main operation to enable interleaving to take place. Without this idling, if $p \neq q$, we have:

$$\begin{aligned} &\text{beh}_{p,\{S\}}.APush(p, x) \wedge \text{beh}_{q,\{S\}}.APush(q, y) \\ \Rightarrow &\quad \text{definitions} \\ &\diamond \mathcal{R}.S.p \wedge \text{OnlyAccessedBy}.S.q \\ \equiv &\quad \text{definitions} \\ &\text{false} \end{aligned}$$

♣

$$\begin{aligned}
APush(p, x) &\hat{=} \text{ENF } OnlyAccessedBy.S.p \bullet SPush(x) \\
AEmpty(p, arv_p) &\hat{=} \text{ENF } OnlyAccessedBy.S.p \bullet SEmpty(arv_p) \\
ADoPop(p, arv_p) &\hat{=} \text{ENF } OnlyAccessedBy.S.p \bullet SPop(arv_p) \\
APop(p, arv_p) &\hat{=} AEmpty(p, arv_p) \sqcap APop(p, arv_p) \\
APP(p) &\hat{=} \text{Idle} ; ((\bigsqcap_{x:Val} APush(p, x)) \sqcap APop(p, arv_p)) ; \text{Idle} \\
AS(P) &\hat{=} \llbracket S / \text{INIT } S = \langle \rangle \bullet \parallel_{p:P} \llbracket arv_p / APP(p)^\omega \rrbracket \rrbracket
\end{aligned}$$

Figure 7: A canonical stack specification




Figure 8: Conflicting execution – interference freedom with two writes




Figure 9: Interference freedom with a read and write

5.3 Behaviour refinement

We prove correctness of the concurrent data structure by proving *refinement* between the concurrent program and an abstract specification and a concrete representation (e.g., in Section 6.3 we show $LS(P)$ is a data refinement of $AS(P)$). The sets of locations of the abstract program may differ from those of the concrete program. Thus, we define a refinement relation between commands parametrised by the sets of abstract and concrete locations and the processes executing the command.

We first consider behaviour refinement — a simple form a refinement in which the concrete and abstract state spaced do not need to be linked.

Definition 5.3 (Behaviour refinement). Suppose A and C are commands in contexts (sets of locations) Y and Z , respectively. We say A is *behaviour refined* by C with respect to a set of processes P (denoted $A \sqsubseteq_P^{Y,Z} C$) iff $beh_{P,Z}.C \Rightarrow beh_{P,Y}.A$ holds.

We write $A \sqsubseteq_P^{Y,Z} C$ for $A \sqsubseteq_{\{p\}}^{Y,Z} C$ (i.e., the set of processes is the singleton set p), write for $A \sqsubseteq_P^Z C$ for $A \sqsubseteq_P^{Z,Z} C$ (i.e., the concrete and abstract contexts are identical) and write $A \sqsubseteq_P C$ for $A \sqsubseteq_P^\emptyset C$ (i.e., the abstract and concrete contexts are empty). Because refinement is defined by universal implication between behaviours, the following monotonicity results may be proved in a straightforward manner using monotonicity of the corresponding interval predicate operators.

Example 5.4. The set of variables of the abstract program is given by:

$$M \hat{=} S \cup \{arv_p \mid p: P\}$$

We perform a behaviour refinement, where we show that the *OnlyAccessedBy* permission in Fig. 7 may be weakened to *IntFree* as given in program $BS(P)$ in Fig. 10, i.e., we prove $AS(P) \sqsubseteq_P^M BS(P)$. Condition $BS(P) \sqsubseteq_P^M AS(P)$ is trivial because $OnlyAccessedBy.S.p \Rightarrow IntFree.S.p$. Condition $BS(P) \sqsubseteq_P^M AS(P)$ holds because for any processes $p, q \in P$ such that $p \neq q$ and intervals Δ_p, Δ_q such that $\Delta_p \cap \Delta_q \neq \emptyset$ (i.e., Δ_p and Δ_q overlap), we have

$$beh_{p,S}.(BDoPush(p, arv_p) \sqcap BDoPop(p, arv_p)).\Delta_p \Rightarrow \neg beh_{q,S}.(BDoPush(q, arv_q) \sqcap BDoPop(q, arv_q)).\Delta_q$$

A visualisation of this is given in Fig. 8, where the write of process p conflicts with the *IntFree.S.q* condition in process q . Note that $AEmpty(p, arv_p) = BEmpty(p, arv_p)$, i.e., one cannot replace the enforced property *OnlyAccessedBy.S.p* by *IntFree.S.p* because for example

$$beh_{p,S}.(\text{ENF } IntFree.S.p \bullet SEmpty(p, arv_p)).\Delta_p \wedge beh_{q,S}.BDoPop(q, arv_q).\Delta_q$$

$$\begin{aligned}
BPush(p, x) &\hat{=} \text{ENF } \text{IntFree}.S.p \bullet SPush(x) \\
BEmpty(p, arv_p) &\hat{=} \text{ENF } \text{OnlyAccessedBy}.S.p \bullet SEmpty(arv_p) \\
BDoPop(p, arv_p) &\hat{=} \text{ENF } \text{IntFree}.S.p \bullet SPop(arv_p) \\
BPop(p, arv_p) &\hat{=} BEmpty(p, arv_p) \sqcap BPop(p, arv_p) \\
BPP(p) &\hat{=} \text{Idle} ; ((\bigsqcap_{x: \text{Val}} BPush(p, x)) \sqcap BPop(p, arv_p)) ; \text{Idle} \\
BS(P) &\hat{=} \llbracket S / \text{INIT } S = \langle \rangle \bullet \parallel_{p:P} \llbracket arv_p / BPP(p)^\omega \rrbracket \rrbracket
\end{aligned}$$

Figure 10: An equivalent specification to Fig. 7

is an allowable behaviour. This may be visualised as shown in Fig. 9, where condition $\text{IntFree}.S.q$ in process q does not conflict with the execution of process q . Replacement of $AS(P)$ by $BS(P)$ simplifies the rest of the proof because the conditions that the implementation needs to ensure are weaker. ♣

Lemma 3. *Suppose P is a non-empty set of processes, $Y, Z \subseteq \text{Var} \cup \text{Addr}$, and A, A_1, A_2, C, C_1 and C_2 are commands such that $A \sqsubseteq_P^{Y,Z} C$, $A_1 \sqsubseteq_P^{Y,Z} C_1$ and $A_2 \sqsubseteq_P^{Y,Z} C_2$ hold, and h and g are interval predicates. Each of the following holds provided that $b' \Rightarrow b, \forall k: \text{Val} \bullet \Diamond(e' = k) \Rightarrow \Diamond(e = k)$ and $h \Rightarrow g$.*

$$[b] \sqsubseteq_P^{Y,Z} [b'] \quad (15)$$

$$v := e \sqsubseteq_P^{Y,Z} v := e' \quad (16)$$

$$A_1 ; A_2 \sqsubseteq_P^{Y,Z} C_1 ; C_2 \quad (17)$$

$$A_1 \sqcap A_2 \sqsubseteq_P^{Y,Z} C_1 \sqcap C_2 \quad (18)$$

$$A^\omega \sqsubseteq_P^{Y,Z} C^\omega \quad (19)$$

$$\text{INIT } g \bullet A \sqsubseteq_P^{Y,Z} \text{INIT } h \bullet C \quad (20)$$

Using an interval-based semantics to formalise a program's behaviour allows one to obtain the following useful results which allows one to split a command executed over a large interval into subintervals, and combine a command executed over a number of adjoining intervals into the same command over the larger interval.

Lemma 4. *For any non-empty set of processes P and a command C , if $\text{beh}_P.C$ joins, then both of the following hold.*

$$C \sqsubseteq_P^{Y,Z} C ; C \quad (21)$$

$$C \sqsubseteq_P^{Y,Z} C ; C^\omega \quad (22)$$

Lemma 5. *For any non-empty set of processes P and a command C , if $\text{beh}_P.C$ splits, then.*

$$C ; C \sqsubseteq_P^{Y,Z} C \quad (23)$$

Note that $C^\omega \sqsubseteq_P^{Y,Z} C$ holds trivially by the definition of $^\omega$. The following lemma states that a guard evaluation is equivalent to a program that performs some finite length idling, some non-empty idling in which c is guaranteed to hold, followed by some more idling.

Lemma 6. $[c] \sqsubseteq_P^Z (\text{ENF } \text{Fin} \bullet \text{Idle}) ; (\text{ENF}(\Box c \wedge \neg \text{Empty} \wedge \text{ReadAllLocs}.c.p) \bullet \text{Idle}) ; \text{Idle}$

Note that the behaviours of both $\text{ENF } \text{Fin} \bullet \text{Idle}$ and Idle hold in any empty interval, and hence each of the properties below follow from Lemma 6:

$$[c] \sqsubseteq_P^Z (\text{ENF}(\Box c \wedge \neg \text{Empty} \wedge \text{ReadAllLocs}.c.p) \bullet \text{Idle}) ; \text{Idle} \quad (24)$$

$$[c] \sqsubseteq_P^Z (\text{ENF } \text{Fin} \bullet \text{Idle}) ; (\text{ENF}(\Box c \wedge \neg \text{Empty} \wedge \text{ReadAllLocs}.c.p) \bullet \text{Idle}) \quad (25)$$

$$[c] \sqsubseteq_P^Z \text{ENF}(\Box c \wedge \neg \text{Empty} \wedge \text{ReadAllLocs}.c.p) \bullet \text{Idle} \quad (26)$$

For example, Lemma 3 and Lemma 6 may be used to prove the following:

$$[b] \sqsubseteq_P^{Y,Z} [b] ; [c] \quad (27)$$

The lemma below allows refinement within a wider context.

Lemma 7. If $A \sqsubseteq_P^{W \cup Y, X \cup Z} C$, $Y \subseteq Z$, $W \cap Y = \emptyset = X \cap Z$ and $W \subseteq X$ then $\llbracket W/A \rrbracket \sqsubseteq_P^{Y,Z} \llbracket X/C \rrbracket$.

Proof. $\begin{aligned} & beh_{P,Z}.\llbracket X/C \rrbracket \\ \equiv & \text{behaviour definition} \\ & Local.Z.p \wedge Z \cap X = \emptyset \wedge beh_{P,Z \cup X}.C \\ \Rightarrow & \text{assumptions} \\ & Local.Y.p \wedge W \cap Y = \emptyset \wedge beh_{P,Y \cup W}.A \\ \equiv & \text{behaviour definition} \\ & beh_{P,Y}.\llbracket W/A \rrbracket \end{aligned}$ \square

We also obtain properties for refinement of enforced conditions. The lemma below states that a refinement may be performed by introducing a new enforced condition or by strengthening an existing enforced condition [17, 19, 22].

Lemma 8. If C is a command, g and h are interval predicates, P is a set of processes and Y, Z are sets of locations, then both of the following hold.

$$C \sqsubseteq_P^{Y,Z} \text{ENF } g \bullet C \quad (28)$$

$$h \Rightarrow g \Rightarrow \text{ENF } g \bullet C \sqsubseteq_P^{Y,Z} \text{ENF } h \bullet C \quad (29)$$

The lemma below allows decomposition within sequential choice and iteration, provided that the interval predicate under consideration joins.

Lemma 9. If C, C_1 and C_2 are commands, g is an interval predicate that joins, Y, Z are sets of locations and P is a set of processes, then

$$\begin{aligned} (\text{ENF } g \bullet C_1 ; C_2) & \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C_1) ; (\text{ENF } g \bullet C_2) \\ (\text{ENF } g \bullet C^\omega) & \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C)^\omega \end{aligned}$$

The following lemma allows decomposition of behaviours with enforced properties over parallel composition, sequential composition and non-deterministic choice. Note that decomposition over sequential composition requires that the interval predicate under consideration splits.

Lemma 10. If $A \hat{=} \parallel_{p:P} A_p$, $C \hat{=} \parallel_{p:P} C_p$, A_1, A_2, C_1 , and C_2 are commands, $Y, Z \subseteq \text{Var} \cup \text{Addr}$ and g and h are interval predicates such that h splits, then both of the following hold.

$$(\forall p: P \bullet A_p \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C_p)) \Rightarrow A \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C) \quad (30)$$

$$(A_1 \sqsubseteq_P^{Y,Z} (\text{ENF } h \bullet C_1)) \wedge (A_2 \sqsubseteq_P^{Y,Z} (\text{ENF } h \bullet C_2)) \Rightarrow (A_1 ; A_2 \sqsubseteq_P^{Y,Z} (\text{ENF } h \bullet C_1 ; C_2)) \quad (31)$$

$$(A_1 \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C_1)) \wedge (A_2 \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C_2)) \Rightarrow (A_1 \sqcap A_2 \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C_1 \sqcap C_2)) \quad (32)$$

6 A coarse-grained linearisable abstraction

In this section, we develop the coarse-grained abstraction of the Trieber stack (Section 6.1) and develop interval-based data refinement rules (Section 6.2). In Section 6.3, we prove that the coarse-grained abstraction is linearisable with respect to the canonical stack specification from Fig. 7 via data refinement and in Section 6.4, we discuss the importance of proving linearisability of coarse-grained abstractions.

6.1 The abstraction $LS(P)$

The coarse-grained abstraction (see Fig. 11) uses Top to obtain a pointer to the top of the stack. In addition, each process uses local variables n_p and rv_p for the new node to be inserted (by the push) and the value to be returned (by the pop), respectively. The stack addresses $SAddr$ are defined to be the set of addresses that are reachable from Top . Hence, for a state σ , we define the following:

$$\begin{aligned} iter_0.\sigma & \hat{=} ptr.(\sigma.Top) \\ iter_{n+1}.\sigma & \hat{=} \begin{cases} null & \text{if } eval.(iter_n.\sigma \mapsto next).\sigma = null \\ eval.(iter_n.\sigma \mapsto next).\sigma & \text{otherwise} \end{cases} \\ SAddr.\sigma & \hat{=} \{node \in Addr \mid \exists n: \mathbb{N} \bullet node \in \{eval.(iter_n \cdot key).\sigma, eval.(iter_n \cdot next).\sigma\}\} \end{aligned}$$

$$\begin{aligned}
EnvSt(p) &\hat{=} ENF \neg WriteSomeLoc.(SAddr \cup \{Top\}).p \bullet \text{True} \\
LSetup(p, x) &\hat{=} newNode(n_p); (n_p \cdot key) := x \\
LDoPush(p) &\hat{=} ENF IntFree.(SAddr \cup \{Top, n_p \cdot key, n_p \cdot next\}).p \bullet \\
&\quad (n_p \cdot next) := ptr.(*Top); Top := (n_p, ctr.(*Top) + 1) \\
LPush(p, x) &\hat{=} LSetup(p, x); EnvSt(p); LDoPush(p) \\
LEmpty(p, rv_p) &\hat{=} [ptr.(*Top) = null]; rv_p := Empty \\
LDoPop(p, rv_p) &\hat{=} ENF IntFree.(SAddr \cup \{Top\}).p \bullet ([ptr.(*Top) \neq null]; rv_p := ptr.(*Top) \mapsto key; \\
&\quad Top := (ptr.(*Top) \mapsto next, ctr.(*Top) + 1)) \\
LPop(p, rv_p) &\hat{=} EnvSt(p); (LEmpty(p, rv_p) \sqcap LDoPop(p, rv_p)) \\
LPP(p) &\hat{=} \text{Idle}; (\prod_{x:Val} LPush(p, x)) \sqcap LPop(p, rv_p); \text{Idle} \\
LS(P) &\hat{=} \llbracket Top, FAddr / \text{INIT } TInit \bullet \parallel_{p:P} \llbracket n_p, rv_p / LPP(p)^\omega \rrbracket \rrbracket
\end{aligned}$$

Figure 11: A coarse-grained abstraction $LS(P)$ of the Treiber Stack

Hence, $SAddr.\sigma$ denotes the addresses that are reachable from Top .

Within program $LS(p)$ in Fig. 11, each process p iteratively chooses executes $LPP(p)$, which at each iteration performs some idling, then for a non-deterministically chosen value x executes $LPush(p, x)$ or $LPop(p, rv_p)$, and then performs some more idling. Within $LPush(p, x)$, command $LSetup(p, x)$ initialises the push and $EnvSt(p)$ allows some execution that does not modify the stack, and $LDoPush(p)$ performs the actual push. The coarse-grained pop is modelled by $LPop(p)$, which allows some initial (non-interfering) idling, then either behaves as $LEmpty(p, rv_p)$, which models a coarse-grained pop on an empty stack, or $LDoPop(p, rv_p)$, which models a coarse-grained pop that removes the top element of a non-empty stack.

Like many CAS-based implementations, Treiber's stack may retry its main operation a number of times before succeeding. An abstraction of this behaviour is modelled by the command $EnvSt(p)$, which guarantees that none of the addresses corresponding to the stack have been modified.

For the rest of this paper, we assume that the set of variables of the coarse-grained abstraction is given by L , which is defined as follows.

$$L \hat{=} \{Top\} \cup FAddr \cup \bigcup_{p:P} \{n_p, rv_p\}$$

Example 6.1. One can prove that the addresses in $SAddr$ (i.e., those corresponding to the stack) are not modified by $LS(P)$. In particular one may prove

$$beh_{P,L}.LS(P) \Rightarrow \forall p:P \bullet \neg WriteSomeLoc.SAddr.p \quad (33)$$

6.2 Data refinement

It is possible to prove behaviour refinement between the coarse-grained program $LS(P)$ and the fine-grained implementation $TS(P)$. However, it is not immediately obvious that $LS(P)$ is linearisable. It turns out that a linearisability proof of $LS(P)$ is non-trivial, however, the proof is simplified because large portion of the code are executed atomically.

As shown by Doherty et al [5, 15, 14] and again by Derrick et al [42, 12, 13], a sound method for proving linearisability is to verify data refinement from an abstract representation of the data structure being implemented to the concrete program. Using the framework of input/output automata [34], Doherty [14] constructs a so-called *canonical automata* in which each operation executes by invoking the operation, then executes an (internal) atomic step (corresponding a step of the data type being implemented) and then returns to an idle state. The argument made is that every trace of the canonical automata is a linearisable [14, 5, 15] and hence any refinement of this automata must also be linearisable. However, because this claim is not formally verified, Derrick et al present an extension that allows links data refinement to the Herlihy and Wing's original definition [42, 13, 12]. In this paper, we apply this extended method to an interval-based setting that allows true concurrency.

An abstract program \mathcal{A} is simulated by concrete program \mathcal{C} with respect to a simulation predicate sim if the initialisation of \mathcal{C} together with sim implies the initialisation of \mathcal{A} and any behaviour of \mathcal{C} over an interval in which sim holds implies the behaviour of \mathcal{A} , and additionally sim holds throughout that interval. For streams s_1 and s_2 , we define

$$s_1 \uplus s_2 \triangleq \lambda t: Time \bullet s_1.t \cup s_2.t$$

If the state spaces corresponding to s_1 and s_2 are disjoint, then for each $t \in Time$, $(s_1 \uplus s_2).t$ is a state and hence $s_1 \uplus s_2$ is a stream.

Definition 6.1 (Data refinement). Suppose $P \subseteq Proc$, $Y, Z \subseteq Var \cup Addr$ are locations such that $Y \cap Z = \emptyset$, $AInit \in StatePred_Y$, and $CInit \in StatePred_Z$. We say $\mathcal{A} \triangleq \llbracket Y/INIT AInit \bullet A \rrbracket$ is *data refined* by $\mathcal{C} \triangleq \llbracket Z/INIT CInit \bullet C \rrbracket$ with respect to a *simulation predicate* $sim \in StatePred_{Y \cup Z}$ iff both:

$$\forall \sigma_a: State_Y, \sigma_c: State_Z \bullet CInit.\sigma_a \wedge sim.(\sigma_a \cup \sigma_c) \Rightarrow AInit.\sigma_a \quad (34)$$

$$\begin{aligned} \forall s_c: Stream_Z, \Delta: Interval, \sigma: State_Y \bullet \\ sim.(\sigma \cup s_c.(glb.\Delta - 1)) \wedge beh_{P,Z}.C.\Delta.s_c \Rightarrow \\ \exists s_a: Stream_Y \bullet (\sigma = s_a.(glb.\Delta - 1)) \wedge \Box sim.\Delta.(s_a \uplus s_c) \wedge beh_{P,Y}.A.\Delta.s_a \end{aligned} \quad (35)$$

This establishes a data refinement [9] between the abstract program A and concrete program C in interval-based setting, where the programs execute in parallel in a truly concurrent manner. As in traditional data refinement, our definition relies on a refinement relation sim which links the concrete and abstract states. Condition (34) ensures that every initialisation of the concrete program that is related to an abstract state via sim must imply an initialisation of the abstract program. By (35), for every concrete stream s_c , interval Δ and abstract state σ , provided that both

1. sim holds between σ and the state of s_c just before Δ and
2. the concrete program executes in s_c over Δ

then there exists an abstract stream s_a such that

1. the state of s_a that immediately precedes Δ is σ ,
2. sim holds in the combined stream $s_a \uplus s_c$ throughout Δ , and
3. the abstract program executes in s_a over Δ .

Proving condition (35) directly is difficult because it does not decompose. However, a predicate of the form $p \Rightarrow (\exists x \bullet q \wedge r)$ holds if both $p \Rightarrow \exists x \bullet q$ and $\forall x \bullet p \wedge q \Rightarrow r$ hold. Hence, (35) can be proved by showing that both of the following hold.

$$\forall s_c: Stream_Z, \Delta: Interval, \sigma: State_Y \bullet \quad (36)$$

$$\begin{aligned} sim.(\sigma \cup s_c.(glb.\Delta - 1)) \wedge beh_{P,Z}.C.\Delta.s_c \Rightarrow \\ \exists s_a: Stream_Y \bullet (\sigma = s_a.(glb.\Delta - 1)) \wedge \Box sim.\Delta.(s_a \uplus s_c) \end{aligned}$$

$$\begin{aligned} \forall s_c: Stream_Z, \Delta: Interval, s_a: Stream_Y \bullet \\ \Box sim.\Delta.(s_a \uplus s_c) \wedge beh_{P,Z}.C.\Delta.s_c \Rightarrow beh_{P,Y}.A.\Delta.s_a \end{aligned} \quad (37)$$

By (36), the simulation condition sim must be such that for any execution of the concrete program that executes from a state that satisfies sim , there must exist an abstract stream such that sim holds throughout the interval. To simplify representation of intervals of the form in (36), we introduce the following notation.

$$\begin{aligned} (g \text{ link}_Y c).\Delta.s &\triangleq \forall \sigma: State_Y \bullet c.(\sigma \cup s.(glb.\Delta - 1)) \wedge g.\Delta.s \Rightarrow \\ &\quad \exists s_y: Stream_Y \bullet \sigma = s_y.(glb.\Delta - 1) \wedge \Box c.\Delta.(s_y \uplus s) \\ g \text{ dref}_{Y,Z} c &\triangleq \forall s_z: Stream_Z, \Delta: Interval \bullet (g \text{ link}_Y c).\Delta.s_z \end{aligned}$$

By (37) for any concrete and abstract streams and interval, if the concrete program executes in the interval and forward simulation holds throughout the interval, then it must be possible to execute the abstract program in the interval and stream. Because $Y \cap Z = \emptyset$, we may further simplify (37) to

$\Box sim \wedge beh_{P,Z}.C \Rightarrow beh_{P,Y}.A$, which may be written using behaviour refinement and enforced properties as

$$A \sqsubseteq_P^{Y,Z} \text{ ENF } \Box sim \bullet C \quad (38)$$

We prove this as follows:

$$\begin{aligned}
& (37) \\
& = Y \cap Z = \emptyset \\
& \quad \forall s_c: Stream_Z, \Delta: Interval, s_a: Stream_Y \bullet \\
& \quad \quad \Box sim.\Delta.(s_a \uplus s_c) \wedge beh_{P,Z}.C.\Delta.(s_a \uplus s_c) \Rightarrow beh_{P,Y}.A.\Delta.(s_a \uplus s_c) \\
& = \text{logic} \\
& \quad \forall s: Stream_{Y \cup Z}, \Delta: Interval \bullet \Box sim.\Delta.s \wedge beh_{P,Z}.C.\Delta.s \Rightarrow beh_{P,Y}.A.\Delta.s \\
& = \text{definition of } '\Rightarrow' \text{ and } '\sqsubseteq' \\
& (38)
\end{aligned}$$

Condition (36) can also be decomposed. First, we prove the following lemma that allows one to decompose a proof of a simulation predicate over chop, which in turn enables decomposition of sequential composition. Similar proof techniques for relational frameworks are well studied [9].

Lemma 11. *Suppose $p \in Proc$, $Y, Z \subseteq Var \cup Addr$ such that $Y \cap Z = \emptyset$, $g_1, g_2 \in IntvPred_Z$ and $c \in StatePred_{Y \cup Z}$. Then $((g_1 ; g_2) \text{ dref}_{Y,Z} c)$ holds if both $(g_1 \text{ dref}_{Y,Z} c)$ and $(g_2 \text{ dref}_{Y,Z} c)$ hold.*

Proof. For an arbitrarily chosen $\sigma \in State_Y$, $\Delta \in Interval$ and $s_z \in Stream_Z$, we have the following calculation.

$$\begin{aligned}
& c.(\sigma \cup s_z.(\text{glb}.\Delta - 1)) \wedge (g_1 ; g_2).\Delta.s_z \\
& = \text{definition of } ';' \\
& \quad c.(\sigma \cup s_z.(\text{glb}.\Delta - 1)) \wedge (\exists \Delta_1, \Delta_2: Interval \bullet (\Delta_1 \cup \Delta_2 = \Delta) \wedge (\Delta_1 \propto \Delta_2) \wedge g_1.\Delta_1 \wedge g_2.\Delta_2.s_z) \\
& \Rightarrow \text{logic and } \text{glb}.\Delta_1 = \text{glb}.\Delta \\
& \quad \exists \Delta_1, \Delta_2: Interval \bullet (\Delta_1 \cup \Delta_2 = \Delta) \wedge (\Delta_1 \propto \Delta_2) \wedge c.(\sigma \cup s_z.(\text{glb}.\Delta - 1)) \wedge g_1.\Delta_1.s_z \wedge g_2.\Delta_2.s_z \\
& \Rightarrow \text{use } (g_1 \text{ dref}_{Y,Z} c) \text{ and logic} \\
& \quad \exists \Delta_1, \Delta_2: Interval, s_y: Stream_Y \bullet (\Delta_1 \cup \Delta_2 = \Delta) \wedge (\Delta_1 \propto \Delta_2) \wedge \\
& \quad \quad \sigma = s_y.(\text{glb}.\Delta_1 - 1) \wedge \Box c.\Delta_1.(s_y \uplus s_z) \wedge g_2.\Delta_2.s_z \\
& = \text{use } \Box c.\Delta_1.(s_y \uplus s_z), \Delta_1 \propto \Delta_2 \\
& \quad \exists \Delta_1, \Delta_2: Interval, s_y: Stream_Y \bullet (\Delta_1 \cup \Delta_2 = \Delta) \wedge (\Delta_1 \propto \Delta_2) \wedge \\
& \quad \quad \sigma = s_y.(\text{glb}.\Delta_1 - 1) \wedge \Box c.\Delta_1.(s_y \uplus s_z) \wedge \\
& \quad \quad (\exists \sigma': State_Y \bullet c.(\sigma' \cup s_z.(\text{glb}.\Delta_2 - 1))) \wedge g_2.\Delta_2.s_z \\
& \Rightarrow \text{logic and } (g_2 \text{ dref}_{Y,Z} c) \\
& \quad \exists \Delta_1, \Delta_2: Interval, s_y, s'_y: Stream_Y, \sigma': State_Y \bullet (\Delta_1 \cup \Delta_2 = \Delta) \wedge (\Delta_1 \propto \Delta_2) \wedge \\
& \quad \quad \sigma = s_y.(\text{glb}.\Delta_1 - 1) \wedge \Box c.\Delta_1.(s_y \uplus s_z) \wedge \\
& \quad \quad \sigma' = s'_y.(\text{glb}.\Delta_2 - 1) \wedge \Box c.\Delta_2.(s'_y \uplus s_z) \\
& \Rightarrow \text{can pick } s''_y \in Stream_Y \text{ such that } \forall t: \Delta_1 \bullet s''_y.t = s_y.t \text{ and } \forall t: \Delta_2 \bullet s''_y.t = s'_y.t, \\
& \quad \text{then rename } s''_y \text{ to } s_y \\
& \quad \exists \Delta_1, \Delta_2: Interval, s_y: Stream_Y, \sigma': State_Y \bullet (\Delta_1 \cup \Delta_2 = \Delta) \wedge (\Delta_1 \propto \Delta_2) \wedge \\
& \quad \quad \sigma = s_y.(\text{glb}.\Delta_1 - 1) \wedge \Box c.\Delta_1.(s_y \uplus s_z) \wedge \\
& \quad \quad \sigma' = s_y.(\text{glb}.\Delta_2 - 1) \wedge \Box c.\Delta_2.(s_y \uplus s_z) \\
& \Rightarrow \text{logic, } \Box c \text{ joins, and } \text{glb}.\Delta_1 = \text{glb}.\Delta \\
& \quad \exists s_y: Stream_Y \bullet \sigma = s_y.(\text{glb}.\Delta - 1) \wedge \Box c.\Delta.(s_y \uplus s_z) \quad \square
\end{aligned}$$

We use the following lemma to further decompose proof obligation (36).

Lemma 12. *If $Y, Z \subseteq Var$ such that $Y \cap Z = \emptyset$, $g \in IntvPred_Z$, $s_z \in Stream_Z$ and $\Delta \in Interval$. Then $(g \text{ link}_Y c).\Delta.s_z$ holds if there exists a $w \in StatePred_Z$ such that:*

$$g.\Delta.s_z \Rightarrow \Box(\Box w \text{ link}_Y c).\Delta.s_z \wedge \Box(\Box \neg w \text{ link}_Y c).\Delta.s_z \quad (39)$$

Proof. For an arbitrarily chosen $\sigma \in State_Y$, $\Delta \in Interval$ and $s_z \in Stream_Z$ and prefix Δ' of Δ , we have that either $\Box w.\Delta'.s_z$ or $\Box \neg w.\Delta'.s_z$ holds. By (39), we have that and a similar calculation to the proof of Lemma 11. \square

$$\begin{aligned}
\text{start_record}_p(v, e) &\triangleq \exists k \bullet \overrightarrow{\text{prev}.e = k} \wedge \overleftarrow{v = k} \wedge \mathcal{W}.v.p \\
\text{end_record}_p(v, e) &\triangleq \exists k \bullet \overrightarrow{e = k} ; [(v = k) \wedge \mathcal{W}.v.p] \\
ARecord(Inv, Res) &\triangleq \text{end_record}_p(HA, HA \cap \langle Inv, Res \rangle) \\
LRecord(Inv, Res) &\triangleq \text{start_record}_p(HL, HL \cap \langle Inv \rangle) \wedge \text{end_record}_p(HL, HL \cap \langle Res \rangle) \\
HAPush(p, x) &\triangleq \text{ENF } ARecord(\text{push}_p^I(x), \text{push}_p^R) \bullet BPush(p, x) \\
HAEEmpty(p, arv_p) &\triangleq \text{ENF } ARecord(\text{pop}_p^I, \text{pop}_p^R(arv_p)) \bullet BEmpty(p, arv_p) \\
HADoPop(p, arv_p) &\triangleq \text{ENF } ARecord(\text{pop}_p^I, \text{pop}_p^R(arv_p)) \bullet BPop(p, arv_p) \\
HAPop(p, arv_p) &\triangleq HAEEmpty(p, arv_p) \sqcap HADoPop(p, arv_p) \\
HAPP(p, arv_p) &\triangleq \text{Idle} ; ((\sqcap_{x:Val} HAPush(p, x)) \sqcap HAPop(p, arv_p)) ; \text{Idle} \\
HAS(P) &\triangleq \llbracket S / \text{INIT } S = \langle \rangle \bullet \parallel_{p:P} HAPP(p)^\omega \rrbracket \\
\\
HLPush(p, x) &\triangleq \text{ENF } LRecord(\text{push}_p^I(x), \text{push}_p^R) \bullet LPush(p, x) \\
HLEEmpty(p, rv_p) &\triangleq \text{ENF } LRecord(\text{pop}_p^I, \text{pop}_p^R(rv_p)) \bullet LEmpty(p, rv_p) \\
HLDoPop(p, rv_p) &\triangleq \text{ENF } LRecord(\text{pop}_p^I, \text{pop}_p^R(rv_p)) \bullet LDoPop(p, rv_p) \\
\\
HLPP(p) &\triangleq \text{Idle} ; ((\sqcap_{x:Val} HLPush(p, x)) \sqcap HLPop(p, rv_p)) ; \text{Idle} \\
HLS(P) &\triangleq \llbracket Top, FAddr / \text{INIT } TInit \bullet \parallel_{p:P} \llbracket n_p, rv_p / HLPP(p)^\omega \rrbracket \rrbracket
\end{aligned}$$

Figure 12: Programs $BS(P)$ and $LS(P)$ extended with history, status and labels



Figure 13: A state corresponding to abstract stack $\langle aa, bb, cc \rangle$, where $ptr.(*Top) = X$ and $ctr.(*Top) = 14$

6.3 Linearisability of $LS(P)$ via simulation

Derrick et al show that a proof of linearisability can be reduced to a proof of data refinement by encoding the definition of linearisability from Section 2.2 within the simulation relation and extending the abstract and concrete programs with histories of invocations and responses [12, 13, 42]. Programs $AS(P)$ and $LS(P)$ extended with histories are respectively given by programs $HAS(P)$ and $HLS(P)$ in Fig. 12. In particular, the canonical program produces a sequential history by recording an invocation immediately followed by the matching response in history HA at the end of the executions of both $HAPush(p, x)$ and $HAPop(p, arv_p)$. On the other hand, the coarse-grained atomic program $LS(P)$ is extended to $HLS(P)$ so that invocations and responses of $HLPush(p, x)$ and $HLPop(p, arv_p)$ in history HL . Note that invocations and responses of operations of $HLS(P)$ may not be sequential, i.e., other processes may be arbitrarily interleaved between any matching pair of events.

In addition, one must also establish a relationship between concrete stack (represented as a linked list) and the abstract stack (represented as a sequence of values). Hence, for an concrete state $\sigma \in State_L$, we define:

$$Stack.\sigma \triangleq \text{let } len = \frac{size.SAddr.\sigma}{2} \text{ in } \lambda n:0..len-1 \bullet (iter_n.Top \mapsto key).\sigma$$

Example 6.2. Consider the abstract stack corresponding to the state σ depicted in Fig. 13. We have

$$\begin{aligned}
iter_0.\sigma &= ptr.(\sigma.Top) & iter_1.\sigma &= eval.(iter_0.\sigma \mapsto next).\sigma \\
&= X & &= eval.(* (X + offset.next)).\sigma \\
& & &= eval.(* (X + 1)).\sigma \\
& & &= \sigma.(X + 1) \\
& & &= Y
\end{aligned}$$

Similarly, $iter_2.\sigma = Z$ and $iter_3.\sigma = null$. Hence, $Stack.\sigma = \langle aa, bb, cc \rangle$

♣

Theorem 6.1. *HAS(P) is data refined by HLS(P) with respect to*

$$simTS \hat{=} (S = Stack) \wedge (\forall p: P \bullet \mathcal{W}.S.p = \mathcal{W}.Top.p) \wedge linearisable(HL, HA) \quad (40)$$

Condition (38) establishes a relationship between concrete stack (which is a linked list) and its abstract representation (which is a sequence), ensures that for any process p , the process has permission to write to S iff it has permission to write to location Top and that that the concrete history HC is linearisable with respect to the abstract history HA .

The linearisation of push and a non-empty pop correspond to the intervals in which Top is modified, and the linearisation point of the empty pop corresponds to an interval in which $ptr.(*Top) = null$ holds. Hence, we use (25) to split the $HLEmpty(p, rv_p)$ operation into the before, during and after cases of the linearisation point as follows.

$$\begin{aligned} HLEmptyPre(p, rv_p) &\hat{=} \text{ENF } \text{start_record}_p(HL, HL \frown \langle pop_p^I \rangle) \wedge \text{Fin} \bullet \text{Idle} \\ HLEmptyLin(p, rv_p) &\hat{=} \text{ENF } \Box(ptr.(*Top) = null) \wedge \neg \text{Empty} \wedge \mathcal{R}.Top.p \bullet \text{Idle} \\ HLEmptyPost(p, rv_p) &\hat{=} \text{ENF } \text{end_record}_p(HL, HL \frown \langle pop_p^I(rv_p) \rangle) \bullet rv_p := \text{Empty} \end{aligned}$$

To prove this, we first show the more straightforward property that for any process p and interval Δ in which p has write permission to Top throughout Δ , p maintains $simTS$ throughout Δ . Furthermore, $simTS$ is also maintained if no process writes to Top .

Lemma 13. *Both of the following hold:*

$$beh_{P,L}.HLS(P) \Rightarrow \Box(\Box(\exists p: P \bullet \mathcal{W}.Top.p) \text{ link}_M simTS) \quad (41)$$

$$beh_{P,L}.HLS(P) \Rightarrow \Box(\Box(\forall p: P \bullet \neg \mathcal{W}.Top.p) \text{ link}_M simTS) \quad (42)$$

41. This holds because there is always a corresponding abstract stack after an update to Top and furthermore intervals in which Top is modified can be treated as linearisation intervals by appending corresponding invocation and response events to the abstract history. \square

42. The only way to modify $simTS$ is to write to HL and the only operations that modify HL without writing to Top are the invocations of each operation and the return operation $HLEmptyPost(p, rv_p)$. Every invocation is trivial because we may treat them as non-pending operations, and $HLEmptyPost(p, rv_p)$ is satisfied by treating operations that have executed $HLEmptyLin(p, rv_p)$ as a pending invocation. \square

We now return to the proof of the theorem that establishes data refinement between $HLS(P)$ and $HAS(P)$.

Theorem 6.1. The proof of the initialisation condition (34) is trivial. The main condition (35) is proved by first splitting the proof into conditions (36) and (37). To prove (36), i.e.,

$$beh_{P,L}.HLS(P) \text{ dref}_{M,L} simTS$$

we apply Lemma 12 with w instantiated to $\exists p: P \bullet \mathcal{W}.Top.p$ and use (41) and (42). We must now prove (37), which we have shown above holds by (38), i.e.,

$$HAS(P) \sqsubseteq_P^{M,L} \text{ENF } \Box simTS \bullet HLS(P)$$

Using Lemma 10 (i.e., decomposition of parallel composition) and Lemma 3 (i.e., monotonicity of ω), this may be proved by showing that for any $p \in P$.

$$HAPP(p) \sqsubseteq_p^{M,L} \text{ENF } \Box simTS \bullet HLPP(p)$$

Because $\text{Idle} \sqsubseteq_P^{Y,Z} \text{Idle}; \text{Idle}$, we may further decompose this using Lemma 10, where the non-idle cases are given below, and the rest of the $HAPP(p)$ are refinements of Idle

$$\begin{aligned} HAPush(p, x) &\sqsubseteq_p^{M,L} \text{ENF } \Box simTS \bullet LDoPush(p) \\ HAEEmpty(p, arv_p) &\sqsubseteq_p^{M,L} \text{ENF } \Box simTS \bullet HLEmptyLin(p, rv_p) \\ HADoPop(p, arv_p) &\sqsubseteq_p^{M,L} \text{ENF } \Box simTS \bullet HLDOPop(p, rv_p) \end{aligned}$$

Each of these proofs is straightforward due to assumption $\Box simTS$. \square

6.4 Remark: The importance of proving linearisability

It is also worth noting that some abstractions of the Treiber Stack cannot be linearised, which reinforces the importance of proving linearisability of a coarse-grained abstraction, as opposed to methods such as [45], which only prove refinement between a concurrent data structure and its coarse-grained abstraction without showing that the abstraction itself is linearisable.

Example 6.3. A pop abstraction

$$\begin{aligned} & [ptr.(*Top) \neq null]; \\ \text{ENF } \Box \neg \mathcal{I}.Top.p \bullet rv_p &:= ptr.(*Top) \mapsto key; \\ Top &:= (ptr.(*Top) \mapsto next, ctr.(*Top) + 1) \end{aligned}$$

cannot be proved linearisable because it is possible for Top to change after the $ptr.(*Top) \neq null$ holds within the guard evaluation $[ptr.(*Top) \neq null]$. Hence, for example, when Top is updated, the stack may already be empty. ♣

Example 6.4. It is not necessary to strengthen $LEmpty(p)$ to

$$\text{ENF } \Box \neg \mathcal{I}.Top.p \bullet [ptr.(*Top) = null]; \quad rv_p := Empty$$

because the value of Top is never used in the latter parts of the code. The guard $[ptr.(*Top) = null]$ is merely used to decide whether or not the code should return. Although this strengthening does provide one with a coarse-grained program that is linearisable, the proof that the coarse-grained program is implemented by the Treiber Stack $TS(P)$ will be more difficult to achieve. ♣

7 Compositional proofs

In this section, we describe how the proof of a command may be decomposed into proofs of the subcomponents. In particular, we present rely conditions in Section 7.1, and decomposition Section 7.2 over parallel composition using rely conditions. In Section 7.3, we present a number of high-level transformation rules specific to CAS-based implementations.

7.1 Rely conditions

We introduce constructs for defining a rely condition, which specifies assumptions about the behaviour of the environment [31]. We note that unlike Jones [31], who assumes rely conditions are relations, we rely conditions are interval predicates, allowing specification of properties over an interval. The behaviour of a command with a rely condition is given by the behaviour of the command in an interval in which the rely condition is assumed to hold. That is, the behaviours of the environment *overlap* [18] with those of the program as opposed to *interleave* [31, 41] with the program.

Definition 7.1. For an interval predicate r and command C , we let $(\text{RELY } r \bullet C)$ denote a command with a *rely condition* r , whose behaviour for any set of processes P and set of locations $Z \subseteq \text{Var} \cup \text{Addr}$ is given by

$$beh_{P,Z}.(\text{RELY } r \bullet C) \quad \hat{=} \quad r \Rightarrow beh_{P,Z}.C$$

Hence, $(\text{RELY } r \bullet C)$ consists of an execution of C under the assumption that r holds. Note that if $\neg r$ holds, then the behaviour of $(\text{RELY } r \bullet C)$ is chaotic, i.e., any behaviour is allowed.

This interpretation of rely/guarantee has been shown to be effective for reasoning in a real-time setting, where conflicting updates by a program and its environment are avoided by ensuring the environment variables are disjoint from the program variables [21, 22]. In this paper, a program and its environment may share a common set of locations, hence we use fractional permission to ensure conflicting accesses do not occur.

Example 7.1. Suppose we want show that for an assignment $x := x + 1$, the final value of x is one greater than its initial value provided that the environment does not modify x . We have

$$\begin{aligned}
& beh_{p,Z}.(\text{RELY } \Box \neg \mathcal{I}.x.p \bullet x := x + 1) \\
\equiv & \text{definition of } beh \\
& \Box \neg \mathcal{I}.x.p \Rightarrow \exists k \bullet beh_{p,Z}.[x + 1 = k] ; \text{update}_{p,Z}(x = k) \\
\Rightarrow & \Box c \text{ splits} \\
& \Box \neg \mathcal{I}.x.p \Rightarrow \exists k \bullet (\Box \neg \mathcal{I}.x.p \Rightarrow beh_{p,Z}.[x + 1 = k]) ; \text{update}_{p,Z}(x = k) \\
\Rightarrow & \text{using } \Box \neg \mathcal{I}.x.p \text{ and by definition } beh_{p,Z}.[x + 1 = k] \Rightarrow \Box \neg \mathcal{W}.x.p \\
& \Box \neg \mathcal{I}.x.p \Rightarrow \exists k \bullet \Box(x + 1 = k) \wedge \neg \text{Empty} ; \Box(x = k) \wedge \neg \text{Empty} \\
\Rightarrow & \Box c \Rightarrow \overleftarrow{c} \\
& \Box \neg \mathcal{I}.x.p \Rightarrow \exists k \bullet \overleftarrow{x + 1 = k} \wedge \neg \text{Empty} \wedge \overrightarrow{x = k} \\
\Rightarrow & \text{logic} \\
& \Box \neg \mathcal{I}.x.p \Rightarrow \exists k \bullet \overleftarrow{x = k} \wedge \neg \text{Empty} \wedge \overrightarrow{x = k + 1}
\end{aligned}$$

♣

7.2 Decomposition using rely conditions

One may develop a number of rules for refining commands with rely conditions.

Lemma 14. *Each of the following holds.*

$$\text{RELY } r \bullet C \sqsubseteq_P^{Y,Z} C \quad (43)$$

$$r \Rightarrow r' \Rightarrow \text{RELY } r \bullet C \sqsubseteq_P^{Y,Z} \text{RELY } r' \bullet C \quad (44)$$

$$r \wedge beh_P.C \Rightarrow beh_P.A \Rightarrow \text{RELY } r \bullet A \sqsubseteq_P^{Y,Z} \text{RELY } r \bullet C \quad (45)$$

$$r \wedge beh_P.C \Rightarrow d \Rightarrow \text{RELY } r \bullet \text{ENF } d \bullet C \sqsubseteq_P^{Y,Z} \text{RELY } r \bullet C \quad (46)$$

Rule (43) allows a rely condition to be removed, (44) allows a rely condition to be weakened and by (45), the refinement holds for the rely condition r on both sides if the behaviour of C implies the behaviour of A under rely condition r . By (46), we may remove the enforced property d if the rely condition and behaviour of C together imply d . Of course, it may be the case that only the rely condition without the program or the program without the rely condition is enough to establish the enforced property. Both these cases are covered by (46).

The lemma below allows one to distribute a rely condition in and out of a sequential composition and an iterated command. The lemma requires that the given rely condition splits.

Lemma 15. *Suppose r is an interval predicate, P a non-empty set of processes and Y, Z are sets of variables. If r splits, then*

$$(\text{RELY } r \bullet C_1 ; C_2) \sqsubseteq_P^{Y,Z} (\text{RELY } r \bullet C_1) ; (\text{RELY } r \bullet C_2) \quad (47)$$

$$\text{RELY } r \bullet C^\omega \sqsubseteq_P^{Y,Z} (\text{RELY } r \bullet C)^\omega \quad (48)$$

$$(\text{RELY } r \bullet A \sqsubseteq_P^{Y,Z} C) \Rightarrow (\text{RELY } r \bullet A^\omega \sqsubseteq_P^{Y,Z} C^\omega) \quad (49)$$

The following theorem shows that rely and enforced conditions form a Galois connection. Namely, command C refines a command A under rely condition r if and only if the command C with enforced condition r refines A .

Theorem 7.1. $(\text{RELY } r \bullet A) \sqsubseteq_P^{Y,Z} C = A \sqsubseteq_P^{Y,Z} (\text{ENF } r \bullet C)$

$$\begin{aligned}
& \text{Proof.} \quad (\text{RELY } r \bullet A) \sqsubseteq_P^{Y,Z} C \\
& = \text{definitions and logic} \\
& \quad r \wedge beh_{P,Z}.C \Rightarrow beh_{P,Y}.A \\
& = \text{definitions} \\
& \quad A \sqsubseteq_P^{Y,Z} (\text{ENF } r \bullet C)
\end{aligned}$$

□

To see the usefulness of this theorem, consider the lemma below that allows refinement over non-deterministic choice in the presence of a rely condition. Using Theorem 7.1, and the result of Lemma 16 (below), we obtain a dual result (50) below on an enforced property.

Lemma 16. $((\text{RELY } r \bullet A_1) \sqsubseteq_P^{Y,Z} C_1) \wedge ((\text{RELY } r \bullet A_2) \sqsubseteq_P^{Y,Z} C_2) \Rightarrow (\text{RELY } r \bullet A_1 \sqcap A_2) \sqsubseteq_P^{Y,Z} C_1 \sqcap C_2$

Using Theorem 7.1, the results (28) and Lemma 16 may both immediately be converted into a property for enforced conditions:

$$\begin{aligned} (A_1 \sqsubseteq_P^{Y,Z} (\text{ENF } r \bullet C_1)) \wedge \\ (A_2 \sqsubseteq_P^{Y,Z} (\text{ENF } r \bullet C_2)) \quad \Rightarrow \quad A_1 \sqcap A_2 \sqsubseteq_P^{Y,Z} (\text{ENF } r \bullet C_1 \sqcap C_2) \end{aligned} \quad (50)$$

One can also use Theorem 7.1 and Lemma 15 to obtain the following dual property, where we assume g is an interval predicate that splits.

$$(\text{ENF } g \bullet C_1) ; (\text{ENF } g \bullet C_2) \sqsubseteq_P^{Y,Z} (\text{ENF } g \bullet C_1 ; C_2) \quad (51)$$

The following theorem may be used to decompose a proof that an parallel composition of an abstract program is refined by the parallel composition of a concrete program in the context of an overall rely condition r .

Theorem 7.2. $(\text{RELY } r \bullet \parallel_{p:P} A_p) \sqsubseteq_P^{Y,Z} (\parallel_{p:P} C_p)$ holds if there exist $P_1, P_2 \subseteq P$ such that $P = P_1 \cup P_2$ and $P_1 \cap P_2 = \emptyset$ and both of the following hold for some interval predicates r_1 and r_2 .

$$(\text{RELY } r \wedge r_1 \bullet \parallel_{p:P_1} A_p) \sqsubseteq_{P_1}^{Y,Z} (\parallel_{p:P_1} C_p) \quad (52)$$

$$(\text{RELY } r \wedge r_2 \bullet \parallel_{p:P_2} A_p) \sqsubseteq_{P_2}^{Y,Z} (\parallel_{p:P_2} C_p) \quad (53)$$

$$r \wedge \text{beh}_{P_2,Z} \cdot (\parallel_{p:P_2} C_p) \Rightarrow r_1 \quad (54)$$

$$r \wedge \text{beh}_{P_1,Z} \cdot (\parallel_{p:P_1} C_p) \Rightarrow r_2 \quad (55)$$

Proof. (52) \wedge (53)

$$\begin{aligned} &= \text{definition of } \sqsubseteq_P^{Y,Z}, \text{ logic } (a \Rightarrow (b \Rightarrow c)) = (a \wedge b \Rightarrow c) \\ &\quad (r \wedge r_1 \wedge \text{beh}_{P_1,Z} \cdot (\parallel_{p:P_1} C_p) \Rightarrow \text{beh}_{P_1,Y} \cdot (\parallel_{p:P_1} A_p)) \wedge \\ &\quad (r \wedge r_2 \wedge \text{beh}_{P_2,Z} \cdot (\parallel_{p:P_2} C_p) \Rightarrow \text{beh}_{P_2,Y} \cdot (\parallel_{p:P_2} A_p)) \\ &\Rightarrow \text{logic} \\ &\quad r \wedge r_1 \wedge \text{beh}_{P_1,Z} \cdot (\parallel_{p:P_1} C_p) \wedge r_2 \wedge \text{beh}_{P_2,Z} \cdot (\parallel_{p:P_2} C_p) \Rightarrow \text{beh}_{P_1,Y} \cdot (\parallel_{p:P_1} A_p) \wedge \text{beh}_{P_2,Y} \cdot (\parallel_{p:P_2} A_p) \\ &\Rightarrow (54) \text{ and } (55) \\ &\quad r \wedge \text{beh}_{P_1,Z} \cdot (\parallel_{p:P_1} C_p) \wedge \text{beh}_{P_2,Z} \cdot (\parallel_{p:P_2} C_p) \Rightarrow \text{beh}_{P_1,Y} \cdot (\parallel_{p:P_1} A_p) \wedge \text{beh}_{P_2,Y} \cdot (\parallel_{p:P_2} A_p) \end{aligned}$$

Hence we have the following calculation:

$$\begin{aligned} &\exists P_1, P_2 \bullet P_1 \cup P_2 = P \wedge (P_1 \cap P_2 = \emptyset) \wedge (52) \wedge (53) \\ &\Rightarrow \text{calculation above and logic} \\ &\quad \exists P_1, P_2 \bullet (P_1 \cup P_2 = P) \wedge (P_1 \cap P_2 = \emptyset) \wedge r \wedge \text{beh}_{P_1,Z} \cdot (\parallel_{p:P_1} C_p) \wedge \text{beh}_{P_2,Z} \cdot (\parallel_{p:P_2} C_p) \Rightarrow \\ &\quad \quad \quad \text{beh}_{P_1,Y} \cdot (\parallel_{p:P_1} A_p) \wedge \text{beh}_{P_2,Y} \cdot (\parallel_{p:P_2} A_p) \\ &= \text{definitions} \\ &\quad (\text{RELY } r \bullet \parallel_{p:P} A_p) \sqsubseteq_P (\parallel_{p:P} C_p) \quad \square \end{aligned}$$

When modelling a lock-free program [8, 13, 48], one assumes that each process repeatedly executes operations of the data structure, and hence the processes of the system only differ in terms of the process ids. For such programs, a proof of the parallel composition may be simplified to as described by the theorem below.

Theorem 7.3. Suppose $p \in \text{Proc}$ and A and C are commands with input parameter p such that $W \subseteq Y$ and $X \subseteq Z$. Then $(\text{RELY } r \bullet \parallel_{p:P} A(p)) \sqsubseteq_P^{Y,Z} (\parallel_{p:P} C(p))$ holds if the following holds for some interval predicate r_1 and some $p \in P$ where $Q \triangleq P \setminus \{p\}$.

$$(\text{RELY } r \wedge r_1 \bullet A(p) \sqsubseteq_p^{Y,Z} C(p)) \quad \wedge \quad (r \wedge \text{beh}_Q \cdot (\parallel_{q:Q} C(q)) \Rightarrow r_1) \quad (56)$$

Proof. Applying Theorem 7.2 and choosing $P_1 = \{p\}$, and $P_2 = Q$, the proof of

$$(\text{RELY } r \bullet \parallel_{p:P} A(p)) \sqsubseteq_P^{Y,Z} (\parallel_{p:P} C(p))$$

decomposes as follows for some interval predicates r_1 and r_2 .

$$(\text{RELY } r \wedge r_1 \bullet A(p) \sqsubseteq_p^{Y,Z} C(p)) \quad \wedge \quad (\text{RELY } r \wedge r_2 \bullet \parallel_{p:Q} A(p) \sqsubseteq_Q^{Y,Z} \parallel_{p:Q} C(p)) \quad (57)$$

$$(r \wedge \text{beh}_Q \cdot (\parallel_{p:Q} C(p)) \Rightarrow r_1) \quad \wedge \quad (r \wedge \text{beh}_p \cdot C(p) \Rightarrow r_2) \quad (58)$$

The first conjuncts of conditions (57) and (58) hold by assumption (56). Choosing $r_2 = \text{true}$, the proof of the second conjunct of (58) is trivial and the proof of the second conjunct of (57) follows by induction on the size of the set of processes Q . In particular, we use the fact that the behaviour of $\|_{p:\{ \}} A(p)$ is true as the base case of the induction. \square

7.3 Transformation rules

In this section we present a number of additional refinement rules that allow one to transform coarse-grained code into code with finer granularity. These proofs are greatly simplified by the fact that we consider interval-based behaviour, which includes consideration of the possible interference from other processes. The theorems we present are developed around the Treiber Stack example. We anticipate that several more can be developed when considering other examples. We further conjecture that such theorems can also be used to derive concurrent data structures via a series of refinements.

The theorem below allows refinement of an enforced property $\Box(e = va)$ to a command C within $\text{RELY } r \bullet va := e ; C ; [e = va]$ provided that C does not write to va and the rely condition r ensures that the ABA problem does not occur on e .

Theorem 7.4. *Suppose p is a process, $Z \subseteq \text{Var} \cup \text{Addr}$, $v \in Z \cap \text{Var}$, e is an expression, C is a command and r is an interval predicate. If both of the following hold*

$$\text{beh}_{p,Z}.C \Rightarrow \Box \neg \mathcal{W}.v.p \quad (59)$$

$$r \Rightarrow \Box \neg \text{ABA}.e \wedge (\forall q: \text{Proc} \setminus \{p\} \bullet \neg \mathcal{W}.v.q) \quad (60)$$

then

$$\text{RELY } r \bullet v := e ; (\text{ENF } \Box(e = v) \bullet C) ; [e = v] \sqsubseteq_p^Z v := e ; C ; [e = v] \quad (61)$$

Proof. Condition (61) is equivalent to

$$r \wedge \text{beh}_{p,Z}.(v := e ; C ; [e = v]) \Rightarrow \text{beh}_{p,Z}.(v := e ; (\text{ENF } \Box(e = v) \bullet C) ; [e = v])$$

By logic, the antecedent of the formula above is equivalent to

$$r \wedge \text{beh}_{p,Z}.(v := e ; (\text{ENF } (\Box(e = v) \vee \Diamond(e \neq v)) \bullet C) ; [e = v])$$

The $\Box(e = v)$ case is trivial. For the $\Diamond(e \neq v)$ case, we have

$$\begin{aligned} & r \wedge \text{beh}_{p,Z}.(v := e ; (\text{ENF } \Diamond(e \neq v) \bullet C) ; [e = v]) \\ \Rightarrow & \quad (59) \text{ and } (60) \text{ together implies } \forall q: P \bullet \Box \neg \mathcal{W}.v.q \\ & r \wedge \text{beh}_{p,Z}.(v := e ; (\text{ENF } (\Diamond(e \neq v) \wedge \text{stable}.v) \bullet C) ; [e = v]) \\ \Rightarrow & \quad (60), \text{beh}_{p,Z}.[e = v] \Rightarrow \Box \neg \mathcal{W}.v.p \text{ and } \mathbf{HC1} \\ & r \wedge \text{beh}_{p,Z}.(v := e ; (\text{ENF } (\Diamond(e \neq v) \wedge \text{stable}.v) \bullet C) ; (\text{ENF } \text{stable}.v \bullet [e = v])) \\ \Rightarrow & \quad \text{beh}_{p,Z} \text{ definition} \\ & r \wedge \exists k \bullet \Diamond(e = k) ; (\Box(v = k) \wedge \neg \text{Empty}) ; (\Diamond(e \neq v) \wedge \text{stable}.v) ; (\text{stable}.v \wedge \Diamond(e = v)) \\ \Rightarrow & \quad \text{using stability of } v \\ & r \wedge \exists k \bullet \Diamond(e = k) ; (\Box(v = k) \wedge \neg \text{Empty}) ; \Diamond(e \neq k) ; \Diamond(e = k) \\ \Rightarrow & \quad \text{using } r \Rightarrow \neg \text{ABA}.e \\ & \text{false} \end{aligned} \quad \square$$

The theorem below allows one to replace an assignment to an expression by an equivalent assignment provided that the process under consideration has the necessary permissions.

Theorem 7.5 (Replace assignment). *Suppose e and e' are expressions, $v \in \text{Var}$, $Y \subseteq Z \subseteq (\text{Var} \cup \text{Addr})$ are sets of locations and p is a process. If $\text{IntFree}.e.p \wedge \text{beh}_{p,Z}.\text{Idle} \Rightarrow \text{ReadAllLocs}.e.p$, then*

$$\text{RELY } \Box(e = e') \wedge \text{IntFree}.e.p \bullet v := e \sqsubseteq_p^{Y,Z} v := e'$$

Proof. The proof holds if $\Box(e = e') \wedge \text{IntFree}.e.p \wedge \text{beh}_{p,Z}.(v := e') \Rightarrow \text{beh}_{p,Y}.(v := e)$

$$\begin{aligned}
& \Box(e = e') \wedge \text{IntFree}.e.p \wedge \text{beh}_{p,Z}.(v := e') \\
\equiv & \text{definition of } \text{beh}_{p,Z} \\
& \Box(e = e') \wedge \text{IntFree}.e.p \wedge \exists k \bullet \text{eval}_{p,Z}.(e', k); \text{update}_{p,Z}(v, k) \\
\equiv & \text{definition of } \text{eval} \\
& \Box(e = e') \wedge \text{IntFree}.e.p \wedge \exists k \bullet (\Diamond(e' = k \wedge \text{ReadAllLocs}.e'.p) \wedge \text{beh}_{p,Z}.\text{Idle}); \text{update}_{p,Z}(v, k) \\
\Rightarrow & \text{using } \Box(e = e') \text{ and assumption } \text{IntFree}.e \wedge \text{beh}_{p,Z}.\text{Idle} \Rightarrow \text{ReadAllLocs}.e.p \\
& \exists k \bullet (\Diamond(e = k \wedge \text{ReadAllLocs}.e.p) \wedge \text{beh}_{p,Z}.\text{Idle}); \text{update}_{p,Z}(v, k) \\
\Rightarrow & Y \subseteq Z, \text{ definition of } \text{beh}_{p,Y} \\
& \text{beh}_{p,Y}.(v := e)
\end{aligned}$$

□

The following theorem allows one to split a guard evaluation so that the variable being tested is stored locally. The theorem allows one to perform some additional behaviour that does not affect the abstract state.

Theorem 7.6 (Introduce command). *Suppose e is an expression, v is a variable, ae is an address-valued expression, p is a process, $Y, Z \subseteq \text{Var}$ such that both $Y \subseteq Z$ and $v \in Y$, and C is a command. Then both of the following hold:*

$$v := e \sqsubseteq_p^{Y,Z} C; v := e \quad \text{provided } \text{beh}_{p,Z}.C \Rightarrow \text{idle}_{p,Y} \quad (62)$$

$$ae := e \sqsubseteq_p^{Y,Z} C; ae := e \quad \text{provided } \text{beh}_{p,Z}.C \Rightarrow \text{idle}_{p,Y} \wedge \neg \text{WriteSomeLoc}.ae.p \quad (63)$$

Proof. We prove (62) as follows. The proof of (63) follows the same structure.

$$\begin{aligned}
& \text{beh}_{p,Z}.(C; v := e) \\
\equiv & \text{definition of } \text{beh}_{p,Z} \\
& \text{beh}_{p,Z}.C; \exists k \bullet \text{eval}_{p,Z}(e, k); \text{update}_{p,Z}(v, k) \\
\Rightarrow & \text{assumption on } \text{beh}_{p,Z}.C \\
& \text{idle}_{p,Y}; \exists k \bullet \text{eval}_{p,Z}(e, k); \text{update}_{p,Z}(v, k) \\
\Rightarrow & \text{logic, } Y \subseteq Z \text{ and } v \in Y \\
& \exists k \bullet \text{idle}_{p,Y}; \text{eval}_{p,Y}(e, k); \text{update}_{p,Y}(v, k) \\
\Rightarrow & (13) \text{ and definition of } \text{beh}_{p,Y} \\
& \text{beh}_{p,Y}.(v := e)
\end{aligned}$$

□

Guard and expression evaluation occur in two different (observable) states, and hence, for example, $[a = 42]; v := a$ does not guarantee that v has a final value 42 because the value of a may change after the guard evaluation. The following theorem shows that testing a guard can be split over multiple steps by introducing a new local variable to the program.

Theorem 7.7 (Split guard). *Suppose $va, v \in Z$, k is a constant and $\text{beh}_{p,Z}.C \Rightarrow \text{idle}_{Y \cup \{va\}}$ then*

$$[v = k] \sqsubseteq_p^{Y,Z} va := v; \text{ENF } \text{IntFree}.va.p \bullet (C; [va = k])$$

Proof. We first prove the following:

$$\begin{aligned}
& \text{IntFree}.va.p \wedge (\text{beh}_{p,Z}.C; \text{eval}_{p,Z}(va, k)) \\
\Rightarrow & \text{assumptions } \text{beh}_{p,Z}.C \Rightarrow \text{idle}_{Y \cup \{va\}} \text{ and } Y \subseteq Z \\
& \text{stable}.va \wedge (\text{idle}_{p,Y}; \text{eval}_{p,Y}(va, k)) \\
\Rightarrow & (13) \\
& \text{stable}.va \wedge \text{eval}_{p,Y}(va, k) \\
\Rightarrow & \text{use } \text{stable}.va \\
& \Box(va = k) \wedge \text{idle}_{p,Y}
\end{aligned}$$

Then, we have the following calculation

$$\begin{aligned}
& \text{beh}_{p,Z}.(va := v; \text{ENF } \text{IntFree}.va.p \bullet (C; [va = k])) \\
\Rightarrow & \text{expand definition} \\
& (\exists j \bullet \text{eval}_{p,Z}(v, j); \text{update}_{p,Z}(va, j)); (\text{IntFree}.va.p \wedge (\text{beh}_{p,Z}.C; \text{eval}_{p,Z}(va, k))) \\
\Rightarrow & \text{calculation above}
\end{aligned}$$

$$\begin{aligned}
& (\exists j \bullet \text{eval}_{p,Z}(v, j) ; \text{update}_{p,Z}(va, j)) ; \Box(va = k) \wedge \text{idle}_{p,Y} \\
\Rightarrow & \text{logic, assume } j \text{ fresh} \\
& \exists j \bullet \text{eval}_{p,Z}(v, j) ; \text{update}_{p,Z}(va, j) ; \Box(va = k) \wedge \text{idle}_{p,Y} \\
\Rightarrow & \text{case analysis on } j = k, \text{ case } j \neq k \text{ yields a contradiction} \\
& \text{eval}_{p,Z}(v, k) ; \text{update}_{p,Z}(va, k) ; \text{idle}_{p,Y} \\
\Rightarrow & Y \subseteq Z, \text{ definition of } \text{update} \text{ using assumption } va \notin Y \\
& \text{eval}_{p,Y}(v, k) ; \text{idle}_{p,Y} ; \text{idle}_{p,Y} \\
\Rightarrow & (14) \text{ twice and definition of } beh \\
& beh_{p,Y}.[v = k]
\end{aligned}$$

□

A theorem such as Theorem 7.7 is more difficult to establish in a model that only considers pre/post states because the guard evaluation on the left of $\sqsubseteq_p^{Y,Z}$ is over a single state and the command on the right is over multiple states.

The following theorem allows an assignment to be introduced provided that the new variable is distinct from the abstract context.

Theorem 7.8. *Suppose $Y, Z \subseteq \text{Var}$ such that $Y \subseteq Z$, $v \in Z \setminus Y$ and e is an expression and $p \in \text{Proc}$. Then*

$$\text{ENF } \neg \text{Empty} \bullet \text{Idle} \sqsubseteq_p^{Y,Z} v := e$$

Proof. The refinement holds because $v \notin Y$, $Y \subseteq Z$ and $v := e$ ensures $\text{idle}_{Z \setminus \{v\}}$. □

We also develop a transformation theorem for executing a successful CAS operation.

Theorem 7.9 (Introduce CAS). *Suppose ae is an address-valued expression, e is an expression, $\alpha \in Z$ and r is an interval predicate such that*

$$r \Rightarrow \Box(\text{accessed}.(*ae) \subseteq Z) \wedge \text{IntFree}.\alpha \wedge \quad (64)$$

$$\Box((\ast ae = \alpha) \Rightarrow (\beta = ae \mapsto f)) \quad (65)$$

$$\Box(\text{ReadAllLocs}.\beta.\alpha \Rightarrow \text{ReadAllLocs}.e) \quad (66)$$

Then

$$\text{RELY } r \bullet ae := e \sqsubseteq_p^Z \text{CASOK}_p(ae, \alpha, \beta)$$

Proof. Using Theorem 7.1, we may equivalently prove $ae := e \sqsubseteq_p^Z \text{ENF } r \bullet \text{CASOK}_p(ae, \alpha, \beta)$. We have the following calculation.

$$\begin{aligned}
& beh_{p,Z}.(\text{ENF } r \bullet \text{CASOK}_p(ae, \alpha, \beta)) \\
\equiv & \text{expand definitions} \\
& r \wedge \text{IntFree}.ae.p \wedge (beh_{p,Z}.[*ae = \alpha] ; beh_{p,Z}.(ae := \beta)) \\
\equiv & \text{expand definitions} \\
& r \wedge \text{OnlyAccessedBy}.(*ae).p \wedge \\
& (\text{eval}_{p,Z}.(*ae = \alpha, \text{true}) ; \exists a, k \bullet (\text{eval}_{p,Z}(ae, a) \wedge \text{eval}_{p,Z}(\beta, k)) ; \text{update}_{p,Z}(a, k)) \\
\Rightarrow & \text{logic, expand definitions} \\
& r \wedge \text{OnlyAccessedBy}.(*ae).p \wedge \\
& \exists a, k \bullet (\Diamond(*ae = \alpha) \wedge \text{idle}_{p,Z}) ; (\text{eval}_{p,Z}(ae, a) \wedge \text{eval}_{p,Z}(\beta, k)) ; \text{update}_{p,Z}(a, k) \\
\Rightarrow & (13) \\
& r \wedge \text{OnlyAccessedBy}.(*ae).p \wedge \\
& \exists a, k \bullet (\Diamond(*ae = \alpha) \wedge \text{eval}_{p,Z}(ae, a) \wedge \text{eval}_{p,Z}(\beta, k)) ; \text{update}_{p,Z}(a, k) \\
\Rightarrow & (64), \alpha \in Z \text{ and } \text{OnlyAccessedBy}.(*ae).p \text{ implies all locations in } *ae = \alpha \text{ are stable} \\
& r \wedge \exists a, k \bullet (\Box(*ae = \alpha) \wedge \text{eval}_{p,Z}(ae, a) \wedge \text{eval}_{p,Z}(\beta, k)) ; \text{update}_{p,Z}(a, k) \\
\Rightarrow & (65) \\
& r \wedge \exists a, k \bullet (\Box(\beta = e) \wedge \text{eval}_{p,Z}(ae, a) \wedge \text{eval}_{p,Z}(\beta, k)) ; \text{update}_{p,Z}(a, k) \\
\Rightarrow & \Box(\beta = e) \text{ and } (66) \\
& r \wedge \exists a, k \bullet (\text{eval}_{p,Z}(ae, a) \wedge \text{eval}_{p,Z}(e, k)) ; \text{update}_{p,Z}(a, k) \\
\Rightarrow & \text{definitions} \\
& beh_{p,Z}.(ae := e)
\end{aligned}$$

□

8 Behaviour refinement the Treiber Stack

In this section, we verify that the Treiber Stack (modelled by $TS(P)$) refines the coarse-grained abstract program (modelled by $LS(P)$), i.e., we prove:

$$LS(P) \sqsubseteq_P TS(P) \quad (67)$$

Our proof strategy decomposes the parallel composition (Section 8.1), which allows us to consider the push and pop operations executed by a single process separately (Section 8.2 and Section 8.3). We derive the necessary rely conditions for the push and pop operations as part of the proofs in Sections 8.2 and 8.3, which are then discharged in Section 8.4.

8.1 Decompose parallel composition

We first apply Lemma 7 and then (20) of Lemma 3 to reduce the refinement to the following proof obligation:

$$\|_{p:P} LP(p) \sqsubseteq_p^{TF} \|_{p:P} TP(p) \quad (68)$$

where

$$\begin{aligned} TF &\triangleq \{Top, FAddr\} \\ LP(p) &\triangleq \llbracket n_p, rv_p / LPP(p)^\omega \rrbracket \\ TP(p) &\triangleq \llbracket t_p, n_p, tn_p, rv_p / TPP(p)^\omega \rrbracket \end{aligned}$$

Using Theorem 7.2, we further decompose the parallel composition in (68) to obtain the following proof obligations.

$$RELY\ r_1 \bullet LPP(p)^\omega \sqsubseteq_p^{TF} PP(p)^\omega \quad (69)$$

$$beh_{P', TF}(\|_{P'} PP(p)^\omega) \Rightarrow r_1 \quad (70)$$

Condition r_1 is yet to be determined, and is calculated as part of the proof of (69). However, we require that the condition r_1 that we derive splits to allow our transformation lemmas to be applied. Recalling that the L denotes the set of variables of the coarse-grained abstraction $LS(P)$, we define

$$T \triangleq L \cup \{p: Proc \bullet tn_p, rv_p\} \quad (71)$$

to be set of variables of the coarse-grained abstraction $TS(P)$ and obtain the follows.

$$\begin{aligned} &(69) \\ \Leftarrow &(48) \text{ decompose iteration assuming } r_1 \text{ splits} \\ &RELY\ r_1 \bullet LPP(p) \sqsubseteq_p^{TF} PP(p) \\ \Leftarrow &(47) \text{ of Lemma 15 assuming } r_1 \text{ splits, then Lemma 16} \\ &RELY\ r_1 \bullet LPush(p) \sqsubseteq_p^{L, T} Push(p) \wedge \quad (72) \\ &RELY\ r_1 \bullet LPop(p) \sqsubseteq_p^{L, T} Pop(p) \quad (73) \end{aligned}$$

8.2 Proof of push operation (72)

We prove the $Push(p)$ operation as follows:

$$\begin{aligned} &RELY\ r_1 \bullet LPush(p) \sqsubseteq_p^{L, T} Push(p) \\ \Leftrightarrow &\text{expand definitions} \\ &RELY\ r_1 \bullet LSetup(p, x); EnvSt(p); LDoPush(p) \sqsubseteq_p^{L, T} Setup(p, x); TryPush(p)^\omega; DoPush(p) \\ \Leftarrow &\text{Lemma 15 assuming } r_1 \text{ splits} \\ &RELY\ r_1 \bullet EnvSt(p); LDoPush(p) \sqsubseteq_p^{L, T} TryPush(p)^\omega; DoPush(p) \\ \Leftarrow &\text{Lemmas 5 and 4, using } EnvSt(p) \text{ both splits and joins} \\ &RELY\ r_1 \bullet EnvSt(p); EnvSt(p); LDoPush(p) \sqsubseteq_p^{L, T} TryPush(p)^\omega; DoPush(p) \end{aligned}$$

Using Lemma 15 (i.e., monotonicity of ‘;’) and the assumption that r_1 splits, the final refinement above holds if both of the following hold.

$$RELY\ r_1 \bullet EnvSt(p) \sqsubseteq_p^{L, T} TryPush(p)^\omega \quad (74)$$

$$RELY\ r_1 \bullet EnvSt(p); LDoPush(p) \sqsubseteq_p^{L, T} DoPush(p) \quad (75)$$

Proof of (74).

$$\begin{aligned}
(74) & \\
\Leftrightarrow & \text{ Lemmas 4 and 5, as } EnvSt(p) \text{ both joins and splits} \\
& \text{RELY } r_1 \bullet EnvSt(p)^\omega \sqsubseteq_p^{L,T} TryPush(p)^\omega \\
& (\text{RELY } r_1 \bullet EnvSt(p))^\omega \sqsubseteq_p^{L,T} TryPush(p)^\omega \\
\Leftarrow & \omega \text{ is monotonic} \\
& \text{RELY } r_1 \bullet EnvSt(p) \sqsubseteq_p^{L,T} TryPush(p) \\
\Leftarrow & \text{ Theorem 7.1} \\
& EnvSt(p) \sqsubseteq_p^{L,T} \text{ENF } r_1 \bullet TryPush(p)
\end{aligned}$$

Commands h_3 and h_5 of $TryPush(p)$ trivially satisfy the requirements on the write permissions within $EnvSt(p)$, and hence satisfy $EnvSt(p)$. For command h_4 , we must ensure that $n_p \cdot next \notin SAddr$, otherwise $\neg WriteSomeLoc.SAddr$ may not hold. Hence, we require that r_1 implies:

$$\Box(pc_p = h_4 \Rightarrow (n_p \cdot next) \notin SAddr) \quad (76)$$

Using assumption (76), the proof of (74) is completed.

Proof of (75). By Theorem 7.1, we may turn any rely condition on the left of $\sqsubseteq_p^{L,T}$ into an enforced property on the right. Hence, condition (75) is equivalent to:

$$EnvSt(p) ; LDoPush(p) \sqsubseteq_p^{L,T} \text{ENF } r_1 \bullet DoPush(p) \quad (77)$$

To prove (77), we first simplify the right hand side of $\sqsubseteq_p^{L,T}$. Because Top includes a modification counter and every update to Top increments this counter, each new value of Top is guaranteed to be different from all previous values and hence, the following is trivially guaranteed:

$$\Box \neg ABA.Top \quad (78)$$

We write $C \sqsupseteq_p^{Y,Z} A$ for $A \sqsubseteq_p^{Y,Z} C$ and perform the following calculation.

$$\begin{aligned}
& \text{ENF } r_1 \bullet DoPush(p) \\
\sqsupseteq_p^T & \text{ expand } DoPush(p), \text{ remove labels} \\
& \text{ENF } r_1 \bullet t_p := *Top ; n_p \cdot next := ptr.t_p ; CASOK_p(Top, t_p, (n_p, ctr.t_p + 1)) \\
\sqsupseteq_p^T & \text{ assumption (78) and Theorem 7.4} \\
& \text{ENF } r_1 \bullet t_p := *Top ; \\
& \quad (\text{ENF } (\Box(\neg \mathcal{I}.Top.p \wedge t_p = *Top)) \bullet n_p \cdot next := ptr.t_p) ; CASOK_p(Top, t_p, (n_p, ctr.t_p + 1)) \\
\sqsupseteq_p^T & \text{ Theorem 7.5 then (29) of Lemma 8} \\
& \text{ENF } r_1 \bullet t_p := *Top ; \\
& \quad (\text{ENF } \Box \neg \mathcal{I}.Top.p \bullet n_p \cdot next := ptr.(*Top)) ; CASOK_p(Top, t_p, (n_p, ctr.t_p + 1)) \\
\sqsupseteq_p^T & \text{ assumption } r_1 \text{ splits, then (51) and Lemma 8 to remove enforced property} \\
& (\text{ENF } r_1 \bullet t_p := *Top) ; \quad (79) \\
& (\text{ENF } r_1 \bullet (\text{ENF } \Box \neg \mathcal{I}.Top.p \bullet n_p \cdot next := ptr.(*Top)) ; CASOK_p(Top, t_p, (n_p, ctr.t_p + 1))) \quad (80)
\end{aligned}$$

Hence, by (17) of Lemma 3, the proof of (75) holds if we prove both of the following:

$$EnvSt(p) \sqsubseteq_p^{L,T} \quad (81)$$

$$LDoPush(p) \sqsubseteq_p^{L,T} \quad (82)$$

The proof of (81) is trivial. For (82), we strengthen r_1 so that it implies that there is no interference on the stack nodes if there is no interference on Top and that there is no interference on addresses $n_p \cdot key$ and $n_p \cdot next$, i.e., we require that r_1 satisfies:

$$\Box(\Box \neg \mathcal{I}.Top.p \Rightarrow IntFree.SAddr.p) \wedge IntFree.\{n_p \cdot key, n_p \cdot next\}.p \quad (83)$$

Thus, we have the following calculation.

$$\begin{aligned}
& \text{(80)} \\
& \sqsubseteq_p^T \quad \text{expand definitions} \\
& \quad \text{ENF } r_1 \bullet (\text{ENF } \Box \neg \mathcal{I}. \text{Top}.p \bullet n_p \cdot \text{next} := \text{ptr}.(*\text{Top})) ; \\
& \quad (\text{ENF } \Box \neg \mathcal{I}. \text{Top}.p \bullet [* \text{Top} = t_p] ; \text{Top} := (n_p, \text{ctr}.t_p + 1)) \\
& \sqsupseteq_p^T \quad \Box c \text{ joins for any state predicate } c \text{ and Lemma 9} \\
& \quad \text{ENF } (r_1 \wedge \Box \neg \mathcal{I}. \text{Top}.p) \bullet n_p \cdot \text{next} := \text{ptr}.(*\text{Top}) ; [* \text{Top} = t_p] ; \text{Top} := (n_p, \text{ctr}.t_p + 1) \\
& \sqsupseteq_p^T \quad \text{(63) of Theorem 7.6 and assumption (83)} \\
& \quad \text{ENF } \text{IntFree}.(\text{SAddr} \cup \{ \text{Top}, n_p \cdot \text{key}, n_p \cdot \text{next} \}).p \bullet (n_p \cdot \text{next} := \text{ptr}.(*\text{Top}) ; \text{Top} := \\
& (n_p, \text{ctr}.(*\text{Top}) + 1) \\
& \sqsubseteq_p^T \quad \text{definition of } L\text{DoPush}(p) \\
& \quad L\text{DoPush}(p)
\end{aligned}$$

8.3 Proof of pop operation (73)

We may decompose this operation as follows:

$$\begin{aligned}
& \text{RELY } r_1 \bullet L\text{Pop}(p) \sqsubseteq_p^{L,T} \text{Pop}(p) \\
& \Leftrightarrow \quad \text{expand definitions} \\
& \quad \text{RELY } r_1 \bullet \text{EnvSt}(p) ; (L\text{Empty}(p) \sqcap L\text{DoPop}(p)) \sqsubseteq_p^{L,T} \text{TryPop}(p)^\omega ; (\text{Empty}(p) \sqcap \text{DoPop}(p)) \\
& \Leftarrow \quad \text{distribute '}', \text{EnvSt}(p) \text{ splits} \\
& \quad \text{RELY } r_1 \bullet \text{EnvSt}(p) \sqsubseteq_p^{L,T} \text{TryPop}(p)^\omega \wedge \tag{84} \\
& \quad \text{RELY } r_1 \bullet L\text{Empty}(p) \sqsubseteq_p^{L,T} \text{Empty}(p) \wedge \tag{85} \\
& \quad \text{RELY } r_1 \bullet \text{EnvSt}(p) ; L\text{DoPop}(p) \sqsubseteq_p^{L,T} \text{DoPop}(p) \tag{86}
\end{aligned}$$

Proof of (84). This property holds in a similar manner to the proof of (74). In particular, the proof holds because $\text{TryPop}(p)$ does not modify any location within $\{ \text{Top}, n_p \cdot \text{val} \} \cup \text{SAddr}$.

Proof of (85). This property is trivial using monotonicity properties and Theorem 7.7.

Proof of (86). We strengthen r_1 so that it implies

$$\Box(\Box(*\text{Top} = t_p \wedge pc_p = lt_7) \Rightarrow \Box(tn_p = \text{ptr}.(*\text{Top}) \mapsto \text{next}) \wedge \text{IntFree}. \text{SAddr}.p) \tag{87}$$

By (87), if the global top value Top matches the local copy t_p , then the global next value $\text{Top}.next$ must be the same as the local copy tn_p , and that there is no interference on the locations within SAddr . Using this condition, we prove (86) as follows.

$$\begin{aligned}
& \text{ENF } r_1 \bullet \text{ToCAS}(p) ; lt_7 : \text{CASOK}_p(\text{Top}, t_p, (tn_p, \text{ctr}.t_p + 1)) \\
& \sqsupseteq_p^T \quad \text{expandin definition of } \text{ToCAS}(p) \\
& \quad \text{ENF } r_1 \bullet t_p := *\text{Top} ; [ptr.t_p \neq \text{null}] ; tn_p := ptr.t_p \mapsto \text{next} ; rv_p := ptr.t_p \mapsto \text{key} ; \\
& \quad lt_7 : \text{CASOK}_p(\text{Top}, t_p, (tn_p, \text{ctr}.t_p + 1)) \\
& \sqsupseteq_p^T \quad \text{Theorem 7.4 and (78)} \\
& \quad \text{ENF } r_1 \bullet t_p := *\text{Top} ; \\
& \quad \left(\begin{array}{l} \text{ENF}(\Box \neg \mathcal{I}. \text{Top}.p \wedge *\text{Top} = t_p) \bullet [ptr.t_p \neq \text{null}] ; \\ \quad tn_p := ptr.t_p \mapsto \text{next} ; \\ \quad rv_p := ptr.t_p \mapsto \text{key} \end{array} \right) ; \\
& \quad lt_7 : \text{CASOK}_p(\text{Top}, t_p, (tn_p, \text{ctr}.t_p + 1)) \\
& \sqsupseteq_p^T \quad \text{use } \Box(*\text{Top} = t_p), \text{ then Lemma 8} \\
& \quad \text{beh}_{p,T}. \text{CASOK}_p(\text{Top}, \alpha, \beta) \Rightarrow \neg \mathcal{W}.t_p.p \\
& \quad \text{ENF } r_1 \bullet t_p := *\text{Top} ; \\
& \quad \left(\begin{array}{l} \text{ENF } \Box \neg \mathcal{I}. \text{Top}.p \bullet [ptr.(*\text{Top}) \neq \text{null}] ; \\ \quad tn_p := ptr.(*\text{Top}) \mapsto \text{next} ; \\ \quad rv_p := ptr.(*\text{Top}) \mapsto \text{key} \end{array} \right) ; \\
& \quad lt_7 : \text{CASOK}_p(\text{Top}, t_p, (tn_p, \text{ctr}.(*\text{Top}) + 1)) \\
& \sqsupseteq_p^T \quad \text{(51) using } r_1 \text{ splits, then weaken enforced property}
\end{aligned}$$

$$\begin{aligned}
& t_p := *Top ; \\
& \left(\begin{array}{l} \text{ENF } \Box \neg \mathcal{I}.Top.p \bullet [ptr.(*Top) \neq null] ; \\ \quad tn_p := ptr.(*Top) \mapsto next ; \\ \quad rv_p := ptr.(*Top) \mapsto key \end{array} \right) ; \\
& \text{ENF } r_1 \bullet lt_7 : CASOK_p(Top, t_p, (tn_p, ctr.(*Top) + 1)) \\
\sqsubseteq_p^T & \quad \text{Theorem 7.9 using (87), } CASOK_p(ae, \alpha, \beta) \Rightarrow IntFree.ae.p \\
& t_p := *Top ; \\
& \text{ENF } \Box \neg \mathcal{I}.Top.p \bullet [ptr.(*Top) \neq null] ; tn_p := ptr.(*Top) \mapsto next ; rv_p := ptr.(*Top) \mapsto key ; \\
& \quad Top := (ptr.(*Top) \mapsto next, ctr.(*Top) + 1)
\end{aligned}$$

Using monotonicity of ‘;’, the proof of (86) reduces to the following proof obligations.

$$EnvSt_p \sqsubseteq_p^{L,T} t_p := *Top \quad (88)$$

$$[ptr.(*Top) \neq null] \sqsubseteq_p^{L,T} [ptr.(*Top) \neq null] ; tn_p := ptr.(*Top) \mapsto next \quad (89)$$

The proof of (88) is trivial because $t_p \notin L$. Property (89) holds as follows:

$$\begin{aligned}
& [ptr.(*Top) \neq null] \\
\sqsubseteq_p^L & \quad \text{definitions} \\
& [ptr.(*Top) \neq null] ; (\text{ENF } (\neg \text{Empty}) \bullet \text{Idle}) \\
\sqsubseteq_p^{L,T} & \quad \text{Theorem 7.8 because } tn_p \notin L \\
& [ptr.(*Top) \neq null] ; tn_p := ptr.(*Top) \mapsto next
\end{aligned}$$

8.4 Proof of (70)

The rely condition r_1 is required to imply (76), (78) and (87). We define the weakest possible condition and obtain

$$r_1 \hat{=} (76) \wedge (78) \wedge (87)$$

To prove (76), we show that the condition below holds:

$$\forall q: P \setminus \{p\} \bullet \Box(pc_q = ht_5 \Rightarrow n_q \neq n_p) \wedge \Box(pc_q = lt_7 \Rightarrow tn_q \neq n_p)$$

which ensures that process q can never insert n_p into the queue. The proof of the formula above relies on the fact that $SAddr \cap FAddr = \emptyset$ is an invariant of $TS(P)$. Invariance of $SAddr \cap FAddr = \emptyset$ is straightforward to verify.

To show that processes $q \neq p$ satisfy (87), we must consider commands executed by process q that either make the antecedent true or falsify the consequent of (87). The counter for Top is only incremented and hence process $q \neq p$ cannot make the antecedent of (87) true. Furthermore, the command in process q that falsifies the consequent (i.e., $CASOK_q(Top, t_q, (tn_q, ctr.t_q + 1))$) also falsifies the antecedent.

9 Conclusions and related work

Methods for verifying linearisability have received a large amount of attention in the last few years. Herlihy and Wing’s original paper use possibilities and Owicki/Gries-style [39] proof outlines, which defines the set of possible abstract data structures that corresponds to each point of interleaving. As we have already mentioned, Doherty et al [5, 15] use a simulation-based method using input/output automata, Vafeiadis et al use a framework that combines separation logic and rely/guarantee reasoning [46, 48] and Derrick et al have developed refinement-based methods [11, 12, 13]. O’Hearn et al develop a method using a so-called hindsight lemma [38] and Jonsson presents a method that uses refinement calculus [32]. A number of tool-based methods have also been developed, but these often place restrictions on the final implementation. For instance, Amit et al present static analysis techniques [1], Burckhardt et al [4] develop a tool for checking whether or not an algorithm is *deterministically linearisable* (so that future behaviour need not be considered) and Vafeiadis has developed a tool that can be used to verify linearisability for such deterministically linearisable algorithms [47]. Verification of linearisability using coarse-grained abstraction has been proposed by Turon and Wand, but they do not show that the

abstraction itself is linearisable [45]. Elmas et al [25], and separately Groves [26] use a reduction-based method, but unlike our approach, these methods are not compositional.

Despite this large set of results, due to the complexity of such concurrent data structures a satisfactory scalable solution to verification remains an open problem. The approach proposed by this paper is to split a verification into two phases — the first reduces the size (and hence complexity) of the problem by showing (via a series of small refinements) that the atomicity of an implementation can be made more coarse, leaving one with a simpler program that can be verified to be linearisable. Note we have presented the verification in a different order, i.e., shown linearisability of the abstraction first.

This paper presents a compositional interval-based method of verifying linearisability that does not require one to identify the linearisation points within the concrete code. Instead, we prove that the concrete code implements a coarse-grained abstraction. Due to this coarse granularity, the linearisation points are easier to identify and the proof itself is simpler. Rely/guarantee-style rules together with splits/joins properties are used to develop transformation theorems, which are in turn used to decompose proof obligations. By using an interval-based framework together with fractional permissions we are able to model true concurrency between parallel processes. This also enables reasoning at a finer level of atomicity than is often allowed because we allow reasoning at the level of variable and memory accesses during expression evaluation.

As Bäumler et al point out, reasoning over interval allows one to determine the future behaviour of a program, which in turn allows one to sometimes avoid backwards reasoning [2], e.g., the Michael and Scott queue [36]. Our experiments indicate that interval-based reasoning via coarse-grained abstraction also simplifies proofs of Heller et al’s coarse grained lazy set algorithm [29], which is known to have linearisation points outside the operations being verified [48, 8, 13]. In the terminology of Burckhardt et al, this corresponds to a non-deterministically linearisable program, and hence lies outside the scope of the tools in [1, 4, 47]. We believe that interval-based reasoning allows generality beyond pre/post state reasoning, and that such generalisations are necessary for taming the increasing concurrency in everyday applications.

The methods we have presented have not yet been mechanised and this remains the next obvious extension to this work. We conjecture that the refinement-based framework will also be useful for a derivation, which we aim to explore as part of future work. Such work would draw on, for instance, the derivational approach proposed by Vechev and Yahav [49].

Acknowledgements. Brijesh Dongol and John Derrick are sponsored by EPSRC Grant EP/J003727/1. We thank Lindsay Groves and Ian J. Hayes for their helpful comments on an earlier draft. This paper has benefited from the input of anonymous reviewers.

References

- [1] D. Amit, N. Rinetzky, T. W. Reps, M. Sagiv, and E. Yahav. Comparison under abstraction for verifying linearizability. In Werner Damm and Holger Hermanns, editors, *CAV*, volume 4590 of *Lecture Notes in Computer Science*, pages 477–490. Springer, 2007.
- [2] S. Bäumler, G. Schellhorn, B. Tofan, and W. Reif. Proving linearizability with temporal logic. *Formal Asp. Comput.*, 23(1):91–112, 2011.
- [3] J. Boyland. Checking interference with fractional permissions. In R. Cousot, editor, *SAS*, volume 2694 of *LNCS*, pages 55–72. Springer, 2003.
- [4] S. Burckhardt, C. Dern, M. Musuvathi, and R. Tan. Line-up: a complete and automatic linearizability checker. In B. G. Zorn and A. Aiken, editors, *PLDI*, pages 330–340. ACM, 2010.
- [5] R. Colvin, S. Doherty, and L. Groves. Verifying concurrent data structures by simulation. *Electr. Notes Theor. Comput. Sci.*, 137(2):93–110, 2005.
- [6] R. Colvin and B. Dongol. Verifying lock-freedom using well-founded orders. In C. B. Jones, Z. Liu, and J. Woodcock, editors, *ICTAC*, volume 4711 of *LNCS*, pages 124–138. Springer, 2007.
- [7] R. Colvin and B. Dongol. A general technique for proving lock-freedom. *Sci. Comput. Program.*, 74(3):143–165, 2009.

- [8] R. Colvin, L. Groves, V. Luchangco, and M. Moir. Formal verification of a lazy concurrent list-based set algorithm. In T. Ball and R. B. Jones, editors, *CAV*, volume 4144 of *LNCS*, pages 475–488. Springer, 2006.
- [9] W. de Roever and K. Engelhardt. *Data Refinement: Model-Oriented Proof Methods and their Comparison*, volume 47 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1998.
- [10] W. P. de Roever, F. de Boer, U. Hannemann, J. Hooman, Y. Lakhnech, M. Poel, and J. Zwiers. *Concurrency Verification: Introduction to Compositional and Noncompositional Methods*. Cambridge University Press, 2001.
- [11] J. Derrick, G. Schellhorn, and H. Wehrheim. Proving linearizability via non-atomic refinement. In J. Davies and J. Gibbons, editors, *iFM*, volume 4591 of *LNCS*, pages 195–214. Springer, 2007.
- [12] J. Derrick, G. Schellhorn, and H. Wehrheim. Mechanically verified proof obligations for linearizability. *ACM Trans. Program. Lang. Syst.*, 33(1):4, 2011.
- [13] J. Derrick, G. Schellhorn, and H. Wehrheim. Verifying linearisability with potential linearisation points. In M. Butler and W. Schulte, editors, *FM*, volume 6664 of *LNCS*, pages 323–337. Springer, 2011.
- [14] S. Doherty. Modelling and verifying non-blocking algorithms that use dynamically allocated memory. Master’s thesis, Victoria University of Wellington, 2003.
- [15] S. Doherty, L. Groves, V. Luchangco, and M. Moir. Formal verification of a practical lock-free queue algorithm. In D. de Frutos-Escrig and M. Núñez, editors, *FORTE*, volume 3235 of *LNCS*, pages 97–114. Springer, 2004.
- [16] B. Dongol. Formalising progress properties of non-blocking programs. In Z. Liu and J. He, editors, *ICFEM*, volume 4260 of *LNCS*, pages 284–303. Springer, 2006.
- [17] B. Dongol. *Progress-based verification and derivation of concurrent programs*. PhD thesis, The University of Queensland, 2009.
- [18] B. Dongol, J. Derrick, and I. J. Hayes. Fractional permissions and non-deterministic evaluators in interval temporal logic. In *AVoCS*, 2012.
- [19] B. Dongol and I. J. Hayes. Enforcing safety and progress properties: An approach to concurrent program derivation. In *Australian Software Engineering Conference*, pages 3–12. IEEE Computer Society, 2009.
- [20] B. Dongol and I. J. Hayes. Approximating idealised real-time specifications using time bands. In *AVoCS 2011*, volume 46 of *ECEASST*, pages 1–16. EASST, 2012.
- [21] B. Dongol and I. J. Hayes. Deriving real-time action systems controllers from multiscale system specifications. In J. Gibbons and P. Nogueira, editors, *MPC*, volume 7342 of *LNCS*, pages 102–131. Springer, 2012.
- [22] B. Dongol and I. J. Hayes. Deriving real-time action systems in a sampling logic. *Science of Computer Programming*, 2012. Accepted 17 Oct, 2011.
- [23] B. Dongol and I. J. Hayes. Rely/guarantee reasoning for teleo-reactive programs over multiple time bands. In J. Derrick, S. Gnesi, D. Latella, and H. Treharne, editors, *IFM*, volume 7321 of *LNCS*, pages 39–53. Springer, 2012.
- [24] B. Dongol, I. J. Hayes, L. Meinicke, and K. Solin. Towards an algebra for real-time programs. In W. Kahl and T. G. Griffin, editors, *RAMICS*, volume 7560 of *LNCS*, pages 50–65. Springer, 2012.
- [25] T. Elmas, S. Qadeer, A. Sezgin, O. Subasi, and S. Tasiran. Simplifying linearizability proofs with reduction and abstraction. In J. Esparza and R. Majumdar, editors, *TACAS*, volume 6015 of *LNCS*, pages 296–311. Springer, 2010.

- [26] L. Groves. Verifying Michael and Scott’s lock-free queue algorithm using trace reduction. In J. Harland and P. Manyem, editors, *CATS*, volume 77 of *CRPIT*, pages 133–142, 2008.
- [27] I. J. Hayes. Towards reasoning about teleo-reactive programs for robust real-time systems. In *SERENE ’08*, pages 87–94, New York, NY, USA, 2008. ACM.
- [28] I. J. Hayes, A. Burns, B. Dongol, and C. B. Jones. Comparing models of nondeterministic expression evaluation. Technical Report CS-TR-1273, Newcastle University, 2011.
- [29] S. Heller, M. Herlihy, V. Luchangco, M. Moir, W. N. Scherer III, and N. Shavit. A lazy concurrent list-based set algorithm. *Parallel Processing Letters*, 17(4):411–424, 2007.
- [30] M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.
- [31] C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM Trans. Prog. Lang. and Syst.*, 5(4):596–619, 1983.
- [32] B. Jonsson. Using refinement calculus techniques to prove linearizability. *Formal Asp. Comput.*, 24(4-6):537–554, 2012.
- [33] R. J. Lipton. Reduction: a method of proving properties of parallel programs. *Commun. ACM*, 18(12):717–721, 1975.
- [34] N. Lynch and M. Tuttle. An introduction to input/output automata. *CWI-Quarterly*, 2(3):219–246, 1989.
- [35] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [36] M. M. Michael and M. L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In *The 15th Annual ACM Symposium on Principles of Distributed Computing*, pages 267–275, May 1996.
- [37] B. C. Moszkowski. A complete axiomatization of Interval Temporal Logic with infinite time. In *LICS*, pages 241–252, 2000.
- [38] P. W. O’Hearn, N. Rinetzky, M. T. Vechev, E. Yahav, and G. Yorsh. Verifying linearizability with hindsight. In A. W. Richa and R. Guerraoui, editors, *PODC*, pages 85–94. ACM, 2010.
- [39] S. Owicki and D. Gries. Verifying properties of parallel programs: An axiomatic approach. *Commun. ACM*, 19(5):279–285, 1976.
- [40] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, pages 55–74. IEEE Computer Society, 2002.
- [41] G. Schellhorn, B. Tofan, G. Ernst, and W. Reif. Interleaved programs and rely-guarantee reasoning with ITL. *TIME*, 0:99–106, 2011.
- [42] G. Schellhorn, H. Wehrheim, and J. Derrick. How to prove algorithms linearisable. In P. Madhusudan and S. A. Seshia, editors, *CAV*, volume 7358 of *LNCS*, pages 243–259. Springer, 2012.
- [43] N. Shavit. Data structures in the multicore age. *Commun. ACM*, 54(3):76–84, 2011.
- [44] R. K. Treiber. Systems programming: Coping with parallelism. Technical Report RJ 5118, IBM Almaden Res. Ctr., 1986.
- [45] A. J. Turon and M. Wand. A separation logic for refining concurrent objects. In T. Ball and M. Sagiv, editors, *POPL*, pages 247–258. ACM, 2011.
- [46] V. Vafeiadis. *Modular fine-grained concurrency verification*. PhD thesis, University of Cambridge, 2007.
- [47] V. Vafeiadis. Automatically proving linearizability. In T. Touili, B. Cook, and P. Jackson, editors, *CAV*, volume 6174 of *LNCS*, pages 450–464. Springer, 2010.

- [48] V. Vafeiadis, M. Herlihy, T. Hoare, and M. Shapiro. Proving correctness of highly-concurrent linearisable objects. In J. Torrellas and S. Chatterjee, editors, *PPOPP*, pages 129–136. ACM, 2006.
- [49] M. T. Vechev and E. Yahav. Deriving linearizable fine-grained concurrent objects. In R. Gupta and S. P. Amarasinghe, editors, *PLDI*, pages 125–135. ACM, 2008.